

Free DVD  
7-in-1

# Linux

THE COMPLETE MAGAZINE ON OPEN SOURCE

VOLUME: 06 ISSUE: 08 OCTOBER 2008 116 PAGES ISSUE# 69

For You

# Virtualisation

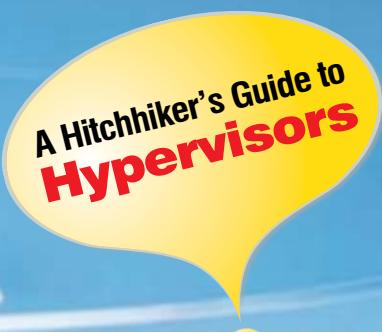
## OpenSolaris

It's ZFS, Zones, Crossbow,  
VirtualBox and Xen

## Virtual Appliances

The Story of Being  
'Just Enough'

## Roll Out a Virtual Infrastructure



A Hitchhiker's Guide to  
**Hypervisors**



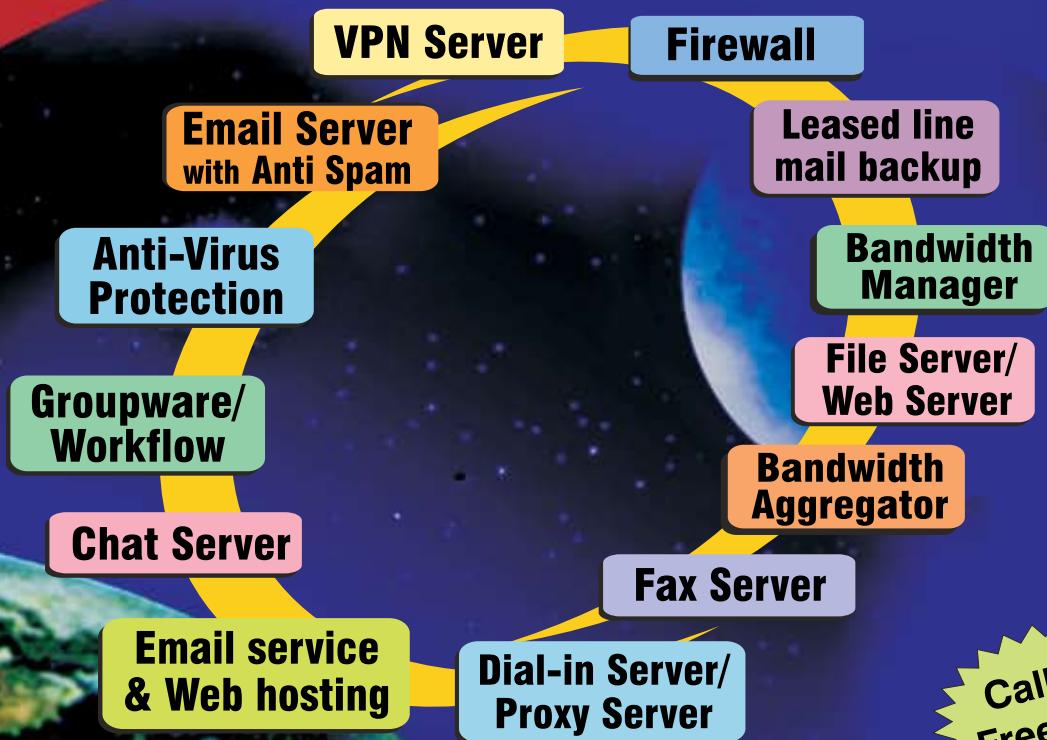
Published by EFY—ISO 9001:2000 Certified

Compiling from Source  
and Kernel Related  
Tweaks

Spying Eyes Around?  
Time to Lock  
Your Data



# Feature packed, complete intranet solution for the discerning enterprise/government organizations/educational and research institutions.



Call for  
Free Trial

- Single SIGN ON system
- GUI based Configuration and administrative program
- Enterprise Back up and Restore functions
- Works with all kinds of Connectivity
- Multi-location support
- Migration from any existing legacy solutions
- Optional Mail/Web hosting services
- 'Optional Clustering Suite' for High availability requirement.

**Carizen**  
SOFTWARE (P) LTD

Chennai : 044-24958222/8228/9296/0610  
Mumbai : 022-66492136/37/38  
New Delhi : 011 - 26239633/4/5  
Bangalore : 080-22208778/9343718998  
Hyderabad : 040-27551584/9391370082

Email: [info@carizen.com](mailto:info@carizen.com)  
Web: [www.carizen.com](http://www.carizen.com)



**Rainmail Intranet Server**

World class network infrastructure at your fingertips

# Linux Learning Centre

*Pioneers in Training on Linux Technologies*

Trained participants from over 36 Countries in 6 Continents

## Linux OS Administration & Security Courses for Migration

- LLC102: Essentials of Linux OS
- LLC103: Linux System & Network Administration
- LLC203: Linux Advanced Administration
- LLC303: Linux System & Network Monitoring Tools
- LLC403: Qmail Server Administration
- LLC404: Postfix Server Administration
- LLC405: Linux Firewall Solutions
- LLC406: OpenLDAP Server Administration
- LLC408: Samba Server Administration
- LLC409: DNS Administration

## Courses for Developers

- LLC104: Linux Internals & Programming Essentials
- LLC105: Programming with Qt
- LLC106: Device Driver Programming on Linux
- LLC107: Network Programming on Linux
- LLC108: Bash Shell Scripting Essentials
- LLC109: CVS on Linux
- LLC204: MySQL on Linux
- LLC205: Programming with PHP
- LLC206: Programming with Perl
- LLC207: Programming with Python
- LLC208: PostgreSQL on Linux
- LLC501: Programming with OpenGL

## RHCE Certification Training

- RH033: Red hat Linux Essentials
- RH133: Red Hat Linux System Administration
- RH253: Red Hat Linux Networking & Security Administration
- RH300/301: Red Hat Rapid Track Certification Course

## RHCSS Certification Training

- RHS333: Red Hat Enterprise Security: Network Services
- RH423: Red Hat Enterprise Directory Services & Authentication
- RHS429: Red Hat Enterprise SELinux Policy Administration

Announcing New LLC  
Technology Series Courses



**Red Hat Fast Track Certification  
Course from 13 October 2008**

**RHCE Exams in Bangalore  
on 14 & 31 October 2008**

**RH423 from 11 Oct '08  
RHS429 from 25 Oct '08  
RHS333 from 04 Oct '08**

**Linux Support & Solutions**  
Installation, Setup and Support Solutions  
for RedHat, Ubuntu, SUSE, CentOS

For more info log on to:  
[www.linuxlearningcentre.com](http://www.linuxlearningcentre.com)  
Call: 9845057731 / 9343780054  
Email: [info@linuxlearningcentre.com](mailto:info@linuxlearningcentre.com)

RED HAT Training Partner  
RHCE Examination Centre



## Linux Learning Centre Private Limited

Registered Office & Corporate Training Centre

# 635, 6th Main Road, (Adj.. Bank of India) Hanumanthnagar, Bangalore 560019

Tel: +91.80.22428538 / 26600839 / 26610999 TelFax: +91.80.26600839

Cell: 9845057731 / 9343780054 Email: [info@linuxlearningcentre.com](mailto:info@linuxlearningcentre.com)

URL: [www.linuxlearningcentre.com](http://www.linuxlearningcentre.com)

LLC Satellite Centre - Bangalore

# 1291, 24th Cross, 30th Main,

BSK II Stage, Bangalore 560070

Tel: +91.80.26712928

# CONTENTS

● OCTOBER 2008

● VOL. 06 NO. 08



ISSN 0974-1054

## Virtualisation

**A Hitchhiker's Guide to Hypervisors — 30**

**Virtual Appliances: It's 'Just Enough' — 35**

**Setting Up A Virtual Infrastructure — 38**

**Virtualisation in OpenSolaris, Part 1: Zones and ZFS — 46**

**Virtualisation in OpenSolaris, Part 2: Crossbow, VirtualBox and Xen — 51**

### FOR YOU & ME

- 18 Miro: Amarok for Video
- 22 FastMail: For Those Who Hate To Wait!
- 26 Software Freedom Day: The Delhi Chapter!
- 106 Open Source: A Weapon of Mass Collaboration

### PLAYERS

- 96 In Pursuit of Freedom!

### DEVELOPERS

- 87 Makefiles for Your Programs

### GEEKS

- 58 Spying Eyes Around? Time to Lock Your Data
- 66 Compiling From Source and Kernel Related Tweaks
- 74 The Building Blocks of Your Drupal Website
- 80 Qingy: An Alternate Desktop Manager
- 84 Programming in Python for Friends and Relations: The Glue for IT
- 90 Creating Beautiful Documents Using LaTeX

# CONTENTS

## ADMIN

- 30 A Hitchhiker's Guide to Hypervisors
- 35 Virtual Appliances: It's 'Just Enough'
- 38 Setting Up A Virtual Infrastructure
- 46 Virtualisation in OpenSolaris, Part 1: Zones and ZFS
- 51 Virtualisation in OpenSolaris, Part 2: Crossbow, VirtualBox and Xen

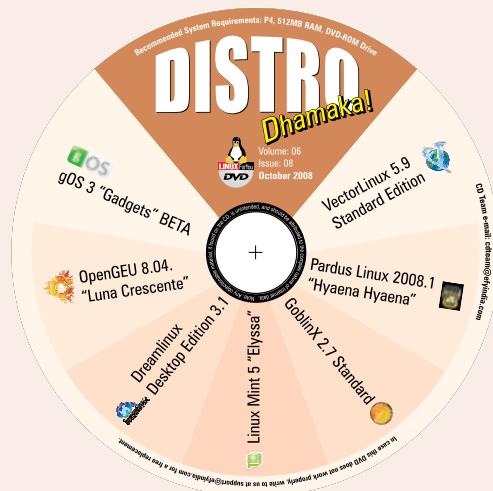
## COLUMNS

- 56 CodeSport
- 83 BraveGNUIndia: Losing That Battle?
- 101 The Joy of Programming: Duff's Device and Some Interesting Aspects of Switch
- 102 A Voyage To The Kernel: Day Four
- 110 FreedomYug: The Harappa of the Future

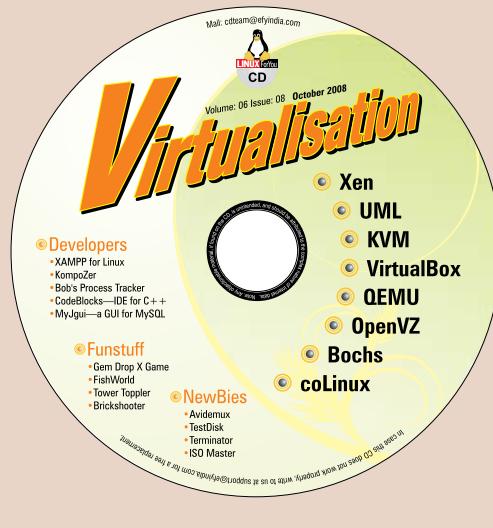
## REGULAR FEATURES

- 06 Editorial
- 08 Feedback
- 10 Technology News
- 16 Q&A Section
- 63 Industry News
- 93 Linux Jobs
- 98 CD Page
- 100 DVD Page
- 108 Tips & Tricks
- 112 FOSS Yellow Pages

## LFY DVD



## LFY CD



# EDITORIAL

Dear Readers,

The BIG news this month for techno-freaks is, of course, the launch of the Google Phone in the US—yes, the rest of the world still has to wait, just like they did for the iPhone. However, us FOSS folks didn't really get too excited over the iPhone, simply because it had literally nothing we could relate to—just another gizmo in a world that already has too many of them! People on the other side of the fence may drawl, “Sour grapes!” Well, so be it! Anyway, is there a reason we *should* care about the new Google Phone?

If you ask me, I guess you should. Why? Because it's Linux-based. But, is it as simple as that? Essentially, the Linux kernel is only one part of Android—a software stack for mobile devices that Google launched a year back. Many FOSS as well as proprietary software, form a part of this software stack that includes an OS, middleware and certain key applications. Considering that a few FOSS tools are also part of the iPhone, how's this one different?

Again, if you ask me, Google's promise to completely open source most of everything by the year end, makes it different. And, that's exactly what makes it BIG news for us. With videos on YouTube and bloggers raving about some of the unique features available on the phone, things are currently looking quite interesting. But we'll only know what it's all about once it's available for public consumption.

Moving on to this month's issue of LFY, from last month's embedded Linux, our focus has shifted to virtualisation in the FOSS world. Starting from a hitchhiker's guide to hypervisor technology and where it's headed, we have tried to address the current state of virtual appliances, and even how to roll out a full-fledged virtual network infrastructure using QEMU and UML. Something interesting that we're able to discuss this year is the state of virtualisation on OpenSolaris—whether it's storage, networks or OS—which we missed out last October.

This month's DVD has something special to offer. We've been getting requests from our readers to go beyond the mainstream distributions to include some that are not so well known. On the same lines, our CD team has managed to pack in seven such distros onto the DVD, each one of which has something unique to offer. And what better time to roll out such a disc than when our issue's focus is on virtualisation? You don't have to install each one of them to try them out—just fire your favourite virtualisation solution and take all of them for a spin.

As I write this editorial, if I were to go back 25 years to September 1983, it was the time RMS conceptualised the GNU project, with an objective to build a completely ‘free’ operating system. Thanks to him and thousands of other developers, we've the luxury of all these FOSS tools now, including more than one full-fledged operating system that has quite a decent mind-share. But how do we grow beyond that? How do we make inroads into the desktop market where another platform dominates with more than 90 per cent of the share? Let me leave you all to ponder on this pressing question for now.

Best wishes,



Rahul Chopra  
Editor, LFY  
rahul@efyindia.com

## Editor

**RAHUL CHOPRA**

## Editorial, Subscriptions & Advertising

### DELHI (HQ)

D-87/1, Okhla Industrial Area,  
Phase I, New Delhi 110020  
Phone: (011) 26810602, 26810603  
Fax: 26817563  
E-mail: info@efyindia.com

### BANGALORE

A-001, Garden Mansions,  
Opp: Anjaneya Complex, Kodihalli,  
Airport Road, Bangalore 560008  
Ph: (080) 25260023; Fax: 25260394  
E-mail: efyblr@efyindia.com

### CHENNAI

M. Nackeeraan  
DBS House, 31A, Cathedral Garden Road  
Near Palmgroove Hotel, Chennai 600034  
Ph: 044-28275191; Mobile: 09962502404  
E-mail: efychn2@efyindia.com

### Customer Care

**E-MAIL:** support@efyindia.com

### Back Issues

#### Kits 'n' Spares

D-88/5, Okhla Industrial Area,  
Phase I, New Delhi 110020  
Phone: (011) 32975879, 26371661-2  
E-mail: kits@efyindia.com  
Website: www.kitsnspares.com

### Advertising

#### KOLKATA

D.C. Mehra  
Ph: (033) 22294788  
Telefax: 22650094  
E-mail: efykal@efyindia.com  
Mobile: 09432422932

### MUMBAI

Flory D'Souza  
Ph: (022) 24950047, 24928520; Fax: 24954278  
E-mail: efyum@efyindia.com

### PUNE

Zakir Shaikh  
Mobile: 09372407753  
E-mail: efpune@efyindia.com

### HYDERABAD

P.S. Muralidharan  
Ph: 09849962660  
E-mail: efyhyd@efyindia.com

### Exclusive News-stand

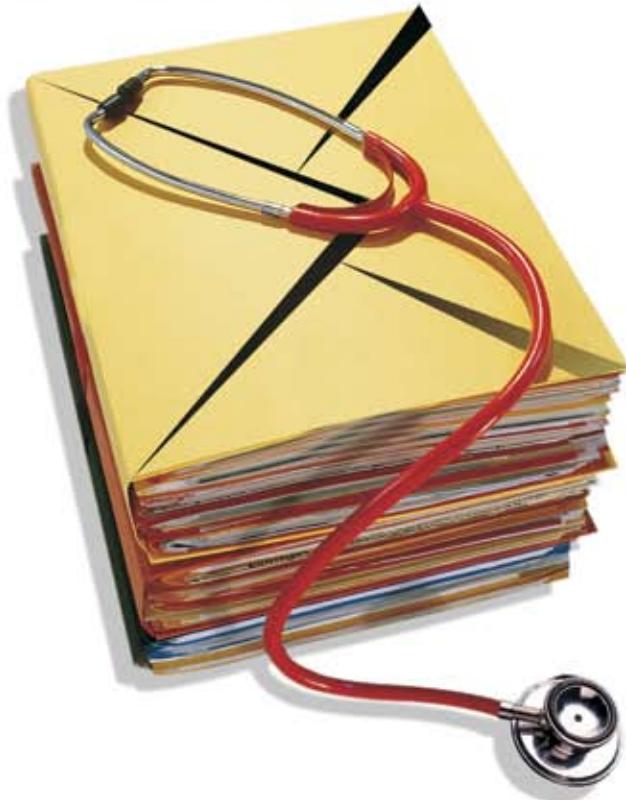
#### Distributor (India)

##### INDIA BOOK HOUSE PVT LTD

Arch No. 30, below Mahalaxmi Bridge, Mahalaxmi,  
Mumbai - 400034 Tel: 24942538, 24925651,  
24927383 Fax: 24950392  
E-mail: info@ibhw.com

Printed, published and owned by Ramesh Chopra. Printed at Ratra Offset, C-101, DDA Shed, Okhla Industrial Area, Phase I, New Delhi 110020, on 28th of the previous month, and published from D-87/1, Okhla Industrial Area, Phase I, New Delhi 110020. Copyright © 2008. All articles in this issue, except for interviews, verbatim quotes, or unless otherwise explicitly mentioned, will be released under Creative Commons Attribution-Share Alike 3.0 Unported License a month after the date of publication. Refer to <http://creativecommons.org/licenses/by-sa/3.0/> for a copy of the licence. Although every effort is made to ensure accuracy, no responsibility whatsoever is taken for any loss due to publishing errors. Articles that cannot be used are returned to the authors if accompanied by a self-addressed and sufficiently stamped envelope. But no responsibility is taken for any loss or delay in returning the material. Disputes, if any, will be settled in a New Delhi court only.

# Ensuring **Healthy** Business Mails - the IT Way



✉ No Virus

✉ No Spam

✉ No Junk Mails

✉ No DoS Attacks

✉ No Phishing Scams

✉ No Directory Attacks

emeric

# CleanM@il™

Securing Business Lifeline



PREMIUM LEVEL  
ANTI-SPAM

**netCORE**  
The Innovation Company

[sales@netcore.co.in](mailto:sales@netcore.co.in) | [www.cleanmail.in](http://www.cleanmail.in)

MUMBAI 9820032324 | NEW DELHI 9350101237 | CHENNAI 9380388685 | BANGALORE 9341609111 | HYDERABAD 9392224001 | PUNE 9373727934 | AHMEDABAD 9824192328 | KOLKATA 9331420404



# You said it...

 I am very thankful to the *LINUX For You* team for presenting us with openSUSE 11.0 with the August 2008 issue. I have a small suggestion—maybe you could provide Slackware Linux 12.1 with your October edition.

—**Sam Benny, by e-mail**

**ED:** *We're glad that you liked openSUSE 11. In fact, we too thought that openSUSE had made a pretty impressive release after a long time. A few of us here have switched to it as our default desktop. Bundling Slackware 12.1 was also on our minds, but due to the overwhelming requests over the months to include distros other than the major ones, we've packed this LFY DVD with ISOs of seven mini distros. Hope you like our selection. Also, as it's been a while since Slackware 12.1 came out, we think it'd be better to wait for a few more months and include their next release instead.*

 I am a computer science lecturer teaching MCA students. I've started teaching Linux now, for the first time. We have Fedora installed on the computers. I would be grateful if you can provide 10 interesting tips on shell programming in *LFY* that I can convert to a lab assignment for the students.

—**Jagannathan Poonkuntran, Coimbatore**

**ED:** *It's great to know your students will be taking up Linux as a part of their curriculum, and hopefully, will continue to use free and open source software in future as well. We also like your suggestion about lab assignments on shell programming. We have included it in our 'to-do' list and will surely discuss it at our team meet.*

 I have been an ardent fan of *LFY*. Though I am not a subscriber, I pick up the news stand copy every month.

I have a collection of the magazines on my bookshelf, from as early as March 2004. The articles covering various Linux distros such as Centos, Debian, Mandrake, and openSUSE have been highly informative. You have even covered some of the UNIX distros such as FreeBSD, OpenSolaris, etc. The quality of the magazine has improved immensely since its inception.

But I have one complaint to make. Correct me if I am wrong, but the one distro you haven't covered till now is IBM AIX. I think it is one of the most rock solid, enterprise-ready products from IBM. IBM certifications are the most sought after in the industry today. Please cover this distro in one of the forthcoming issues of *LFY*. Also, can you tell me of a place in Mumbai where I can find used IBM servers. I want to practise for the AIX exams and need a machine for that. Are there any institutes that provide training on IBM AIX?

—**Mahesh Gurav, Mumbai**

**ED:** *Thank you for all that praise, but it's our job, isn't it? We're glad that you find the content good and useful.*

*There's one problem in covering the AIX platform—it's proprietary. It is based on the original AT&T UNIX code that IBM had licensed in the 80s and then continued to build on top of that. Although we sometimes do cover proprietary solutions, that's only when they run on free platforms, or conversely, when a free solution we're featuring runs on a non-free platform. So, if we talk about different platforms, it's only those that are free—GNU/Linux distros, the BSDs, OpenSolaris, etc. Otherwise, if we talk about a non-free platform like Windows, it's only to highlight that there are free software available for that platform also.*

*But AIX is a completely different ball game. Apart from the fact that it's non-free software, it doesn't run on off-the-shelf hardware too. So, authors*

*who write for us generally do not have much of an idea about how it works.*

 How do I install Knoppix 5.3.1 that came along with the September issue of *LFY*? Knoppix seems to be a more complete system than any other OS that I have ever seen.

—**Nikit Batale, by e-mail**

**ED:** *Knoppix is not meant to be installed on the system. However, Knopper does provide a way to put the distro on the system. Open a shell session and execute the following command and follow the instructions:*

```
sudo knoppix-installer
```

*Note that the installer is command line-based, and makes you use the cfdisk partitioning tool to partition your hard disk. If you don't know how to use it, please consult the cfdisk man page. Also, since it's a Live DVD with more than 12 GB of software, make sure you have a root partition that is more than 13 GB (14 GB recommended). Also, the installer prompts you to choose between three modes of installation: Debian, Knoppix, and Beginner. The project recommends that you go for the Debian option.*

*The DVD has a Knoppix manual with all the details. You can also take a quick look at [www.knoppix.net/wiki/Hd\\_Install\\_HowTo](http://www.knoppix.net/wiki/Hd_Install_HowTo). Also, when the installation starts, the progress bar doesn't move—at least that was the case in our test system. On an Athlon X2 5600+ system with 2 GB of RAM, it took around 40 minutes to install the whole thing. So, on older systems it could easily take more than an hour, considering the amount of data it has to copy to the hard disk—you'll be better off with a newer and faster DVD drive.*

**Please send your comments or suggestions to:**

**The Editor**

**LINUX FOR YOU Magazine**

D-87/1, Okhla Industrial Area, Phase I,

New Delhi 110020; Phone: 011-

26810601/02/03;

Fax: 26817563; Email: [lfyedit@efyindia.com](mailto:lfyedit@efyindia.com);

Website: [www.OpenITis.com](http://www.OpenITis.com)

## India's largest Linux based Enterprise Mailing Solution with 1000 + installations

emeric  
**M@ilServ™**

Flexible. Reliable. Scalable

We thank our valued customers  
for their continued support  
for the last 10 years



Sharekhan



Cipla



*... and many more*

Email : [mailserv@netcore.co.in](mailto:mailserv@netcore.co.in) • Visit us at : <http://www.netcore.co.in/mailserv>

# TECHNOLOGY NEWS

## Fit-PC Slim: Smallest Linux PC

How about a fanless Linux-powered PC that uses a mere 4 to 6 watts of power? CompuLab has launched the Fit-PC Slim, measuring 11x10x3 centimetres, weighing 380 grams, and powered by an AMD Geode LX800 500 MHz processor and 512 MB of RAM. The general specifications of the PC include: VGA output for display, pre-installed 60 GB HDD, 10/100 Ethernet with RJ45 connector, optional 802.11b/g WiFi Ethernet, 3xUSB 2.0 (2 front-panel, 1 rear) and 1x serial (with proprietary connector). The operating temperature of the device is 0 to 45 degrees Celsius.

According to the website, upgrading the hard disk in the Fit-PC Slim is a matter of opening two screws, sliding out the old hard disk and sliding in the new one. Additionally, the Wi-Fi in Fit-PC Slim supports access point mode so the PC can be used as an intelligent wireless router.

Fit-PC Slim Linux is shipped with pre-loaded Ubuntu 8.04 and Gentoo 2008.0 in dual boot mode. For more details about the product, visit [www.fit-pc.com](http://www.fit-pc.com).

## Smallest PC from Dell runs Ubuntu

Computers, which once used to occupy an entire room, today have shrunk to the size of a book. Intel's Atom processor is further pushing the size down, making PCs smaller and smaller. Flowing with the same current, Dell has unveiled the Inspiron Mini 9--a small, easy-to-carry device perfect for surfing the Web, chatting with friends, blogging, streaming content, uploading photos or enjoying favourite online videos, music and games. Dell calls it the best buddy of those who love to stay online.

With a starting weight of 1.3 kilograms, digital nomads will value the Inspiron Mini's durable design, with sealed keyboard and reliable solid state drive (SSD) memory storage. A bright 8.9-inch (22.6 cm) glossy LED display (1024x600) presents most Web pages with no left-right scrolling, and the keypads are large and easy to navigate. Standard built-in Wi-Fi means quick and easy wireless Internet access to hot spots in the home, on campus, in a local coffee shop, in the office or at a conference.

Powered by Intel Atom processor (1.6GHz, 512KB L2 Cache, 533MHz FSB), it runs one of the most popular FOSS operating systems--Ubuntu Linux 8.04--with a custom Dell interface, although users can also opt for Windows XP Home Edition. The Mini 9 can have up to 1GB DDR2 SD RAM, depending on your choice of configuration.

It also has a built-in webcam, bundled with Dell Video Chat, making it easy to stay in touch, recording and sending video e-mails, or even with PC-to-PC phone calls around the world. DVC even supports four-way calling, making virtual family reunions a reality. Its built-in Bluetooth enables easy wireless connections to Bluetooth-enabled accessories stereo headphones, a mouse, a printer, etc. Since the Indian prices are still not available, we recommend keeping your eyes open.



## C-DAC launches advanced version of BOSS

C-DAC has launched its Bharat Operating Systems Solutions (BOSS) Linux software version 3.0, developed by NRCFOSS (National Resource Centre for Free/Open Source Software).



BOSS v3.0 is coupled with GNOME and KDE, and comes with wide Indian language support and packages that are relevant for use in the government domain. The software is also endowed with Bluetooth for short range communications, along with features like a RSS Feed reader and PDF viewer to edit documents.

The ultimate objective of creating BOSS Linux is to enable literate people of India, not conversant with English, to be exposed to the benefits of FOSS and GNU/Linux. With the BOSS v3.0, which has been localised to 18 Indian languages, more people may now turn to FOSS as an alternative to using illegal unlicensed proprietary software. You can point your download managers to [downloads.bosslinux.in/BOSS-3.0/boss-3.0-i386.iso](http://downloads.bosslinux.in/BOSS-3.0/boss-3.0-i386.iso) in order to grab the new version.

## Corel LinDVD now supports ultra-mobile PCs and MIDs

Corel Corporation, a developer of graphics, productivity and digital media software, has announced that Corel LinDVD will now support ultra-mobile PCs (UMPCs) and mobile Internet devices (MIDs), as well as streaming media and a wider range of standard and high-definition video and audio encoding standards. LinDVD is the Linux playback software based on the same industry-leading video technologies that underlie Corel WinDVD, the world's No 1 video and DVD playback software. LinDVD, with support for UMPCs and MIDs, is available to OEMs worldwide.



## Simplified and Scalable Storage Solutions that fit your Budget

### Specification

#### Host Interface:

- SAS : Two Min SAS 4x 3Gb SAS Ports - 300MB/sec per port
- FC : Two 4Gb Fibre Channels - 400MB/sec per channel
- iSCSI : 4 x 1Gbps Ethernet with iSCSI offload

#### Raid Subsystem:

- RAID level 0, 1, 10(1E), 3, 5, 6, 30, 50, 60 or JBOD
- 1GB Battery backed cache with ECC
- Online RAID level/stripe size migration
- Online capacity expansion and RAID level migration simultaneously
- Instant availability and background initialization
- Greater than 2TB per volume set (64-bit LBA support)



#### Drives:

- Supports up to 16-hot-swap SAS/SATAII HDDs in one box
- Supports up to 80 HDDs by cascading multiple JBODs

#### Management:

- LCD Control Panel for setup, alarm mute and configuration
- Firmware-embedded web browser-based RAID manager - access your RAID subsystem from any standard internet browser via 10/100 Lan port

#### Enclosure:

- 3U rack-mountable with rail-kit
- Redundant Power supply

\* Up to 600MB/s throughput and 120,000 IOPS

\* The product picture is representative and the actual product may differ in looks. Please contact us for complete information.

# TECHNOLOGY NEWS

## PC-BSD v 7.0 is now out

The PC-BSD team has announced the availability of PC-BSD version 7.0, codenamed 'Fibonacci'. Marking a milestone for the PC-BSD project by moving to the latest FreeBSD 7-STABLE, the release incorporates the KDE 4.1.1 desktop. According to the announcement: "Users will immediately notice the improved visual interface that KDE 4.1.1 offers, as well as many improvements in hardware support and speed from the update to FreeBSD 7-STABLE. PC-BSD 7.0 also offers a large and growing library of self-contained PBI files available for installation, and improvements for other locales on the PBI Directory web site. This release also offers new methods of installation, including a DVD, USB and Internet / network install." To read the release notes and download the distribution, visit [www.pcbsd.org](http://www.pcbsd.org)



## GIMP 2.5.4 released; v2.6 to follow soon

The GIMP developers have released version 2.5.4 of the popular image manipulation tool, which probably is the last preview version before the big GIMP 2.6.0 release, scheduled for September as we go to press. Changes in GIMP 2.5.4 include: improved look and feel of the navigation dialogue box and navigation pop-ups; improved positioning of the image in the image window; optimised new scaling code; various fixes to the Python bindings; addition of search entry to the keyboard shortcuts and input controller; and a few other handy add-ons. All this is in addition to the general bug fixes and code clean-up.



## Collaborate with GroupWise Open beta

Focused on making employees productive no matter where or how they work, Novell has announced open beta availability of the newest version of Novell GroupWise, a leading collaboration solution. The beta version of Novell GroupWise offers customers and partners a single integrated solution that combines traditional e-mail and calendaring functionality in a personal dashboard with team collaboration workspaces and new Web 2.0 resources, such as wikis, blogs and RSS feeds.

GroupWise offers a wide range of new functionality and collaboration tools. A 'mash-up' style personal productivity dashboard allows users to customise their workspace, providing a comprehensive, single overview of the folders, e-mails, appointments, tasks, collaboration and Web tools that matter most to them. The new contact management features help users track, manage and develop business relationships. GroupWise also supports hundreds of hand-held devices with real-time synchronisation for instant, secure sync of collaboration information between the GroupWise server and the device.

Open beta versions of Novell GroupWise can be downloaded for free at [www.novell.com/groupwisebeta](http://www.novell.com/groupwisebeta)

## LynxOS 5.0 helps create powerful devices using advanced hardware and chipsets

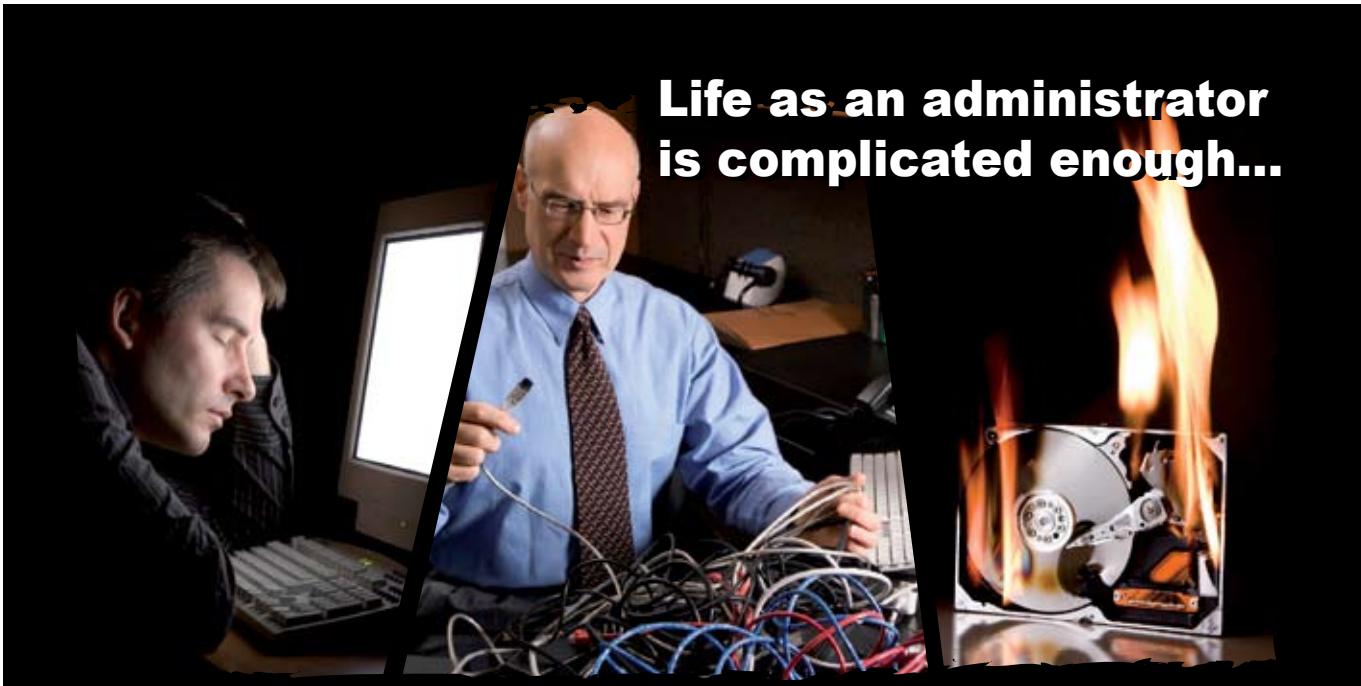
LynxWorks and Realtime Techsolutions have unveiled LynxOS RTOS 5 that will enable software developers to create more powerful devices using advanced hardware and chipsets. One of the key features of LynxOS 5 is its ability to take advantage of the performance gains using symmetric multiprocessing (SMP) architecture. By utilising SMP architecture, in which multi-identical processors are connected in the shared-memory mode, customers using LynxOS 5 will benefit, as the operating system will allow any processor to work on any task, regardless of the position of data in the memory.

Besides the SMP capabilities, because of its rigorous reliability requirements and meticulous adherence to open standards such as POSIX and Linux, LynxOS is in demand in verticals like telecommunications, military/aerospace, industrial, and automotive.

The POSIX interfaces provide advanced real-time and other essential capabilities in the areas of process creation, scheduling, time management, wide characters and dynamic linking. The interfaces also facilitate the migration of legacy POSIX applications along with the creation of new, portable POSIX applications for execution in the LynxOS environment.

LynxOS 5 offers a new Linux application binary interface (ABI) that permits the running of Linux applications along with native POSIX applications, without the need to modify them. This allows customers to leverage several Linux third-party COTS applications. LynxOS 5 also provides advanced networking capabilities.





**Life as an administrator  
is complicated enough...**

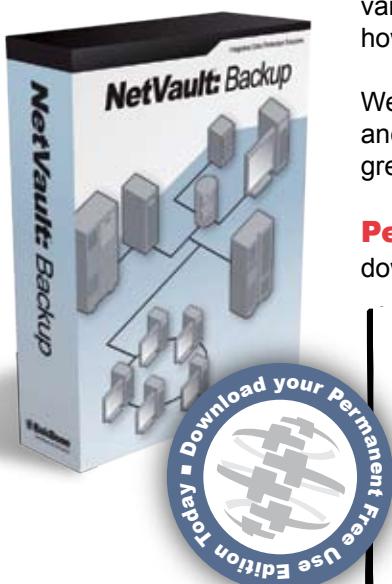
**But Backup and Recovery for your  
Linux Servers does not need to be.**

**NetVault: Backup** simplifies backup and recovery  
without compromising **Functionality** and **Scalability**

NetVault: Backup provides unmatched Data Protection for all major variants of Linux. We are now offering you the chance to see just how good NetVault is at no cost.

We provide continuous data protection (CDP) for your Linux servers and advanced application protection and recovery for MySQL, PostgreSQL, Oracle, Exchange, DB2 to name just a few.

**Permanent FREE Use Edition for Linux** is available for download at <http://www.bakbone.com/nvbu/redhat/>



**NetVault** is a true Linux data protection solution featuring:

- Online backup
- Point and click recovery
- Fully automated protection and recovery support for Linux based applications
- Virtual Tape Library (VTL) with disk staging
- SAN support with LAN free backups.
- Backup duplication for off-site storage of backups
- Support for MySQL native replication, restore to table level and to alternative databases or instances.

# TECHNOLOGY NEWS

## Program in your mother tongue with Hindawi Release 3

The Hindawi Project [[hindawi.in](http://hindawi.in)] has released version 3 of the programming platform, which includes an online edition deployed on Java virtual PC. For those unfamiliar with what it is: Hindawi is a complete non-English systems programming platform that supports all paradigms of programming languages from assembly language to logic and functional programming. It effectively

shatters the language barrier allowing non-English literates to take up computer sciences and participate in the ICT revolution at all levels of technology, in their mother tongue, without the need to master English.

With the new version, you don't need a local installation to get started—you can just go to [hindawi.in/online](http://hindawi.in/online) and learn computing in your mother-tongue—because of being ported to the Java Virtual PC (JPC). The online version can be used on virtually every computer platform, including mobile phones that support Java. You can

use Aadesh (Hindi command shell), Laghu (a simplified Hindawi IDE), and Shaili Robot (Hindi LOGO). This is accompanied by training videos on the top of the page, which will be uploaded on a regular basis. You can view the videos and practice the lessons on the JPC Hindawi screen at the bottom of the page. This method of learning to program attempts to foster self-explanatory and exploratory learning that leads to a deep cognitive understanding of the topic. Students can define their own pace, and re-listen to lectures till they have mastered a particular skill.

Additionally, the new lightweight IDE, Laghu, has been added to Hindawi@DOS. This has mainly been prompted by the limits of the JPC environment. Laghu has two versions: the default supports editing 20 lines of Hindi text, while the Laghu200 version supports 200 lines of Hindi text on JPC and 2,000 lines on the DOSBox and native versions. As Hindawi@DOS is packaged with DOSBox, it allows it to be run on any platform supporting DOSBox, which includes Linux, Windows, Mac OS X, BeOS, BSD and many others.

## Transverse launches blee(p)

Transverse, an open source business solutions company, has launched its open source OSS/BSS platform called Business Logic Execution Environment and Platform, or blee(p). The solution is said to extend the promise of open source computing to telecom operational support systems to leverage better quality, application agility, innovation and lower total cost of ownership.

Transverse said blee(p) takes advantage of the newest technology innovations and the most advanced open source projects to deliver an end-to-end telecom back office that is flexible and adaptable to the rapidly changing needs of carriers. It is designed as a fully integrated set of business management services for back office systems. blee(p) services are grouped into business domain structures that provide more than 2,100 services via meta-domains. Utilising service-oriented architecture (SOA), these domains are easily extended through a plug-in framework, allowing a limitless number of business solutions to be assembled in days, instead of weeks or months.



## Desktop virtualisation with Sun xVM Virtual Box 2.0

Sun Microsystems has announced a new version of Sun xVM VirtualBox, a free and open source desktop virtualisation software, along with 24/7 premium support for enterprise users. With this release, enterprises will be able to fully reap the benefits of the xVM VirtualBox platform and deploy it across their organisations with guaranteed technical support from Sun.

xVM VirtualBox 2.0, released in early September, was soon followed by a bug-fix version 2.0.2 (included in this month's *LFY* CD). The new version comes with support for 64-bit operating systems. The software also offers a new user interface for the Mac platform, improved networking for the Mac OS X and Solaris OS, as well as improved performance, especially on AMD chips. Additionally, customers who purchase an enterprise subscription will also receive a Right-to-Use Licence, allowing them to use the xVM VirtualBox platform with their own software deployment tools.

A mere 20 MB download, xVM VirtualBox software can be installed in less than five minutes. Subscriptions start at \$30 (USD) per user per year, which includes 24/7 support, and discounts are available based on volume. To download the software and sign up for an enterprise support subscription, go to [www.sun.com/software/products/virtualbox/get.js](http://www.sun.com/software/products/virtualbox/get.js). Of course, a user can opt for the open source version (OSE), which is being made available from the official software repositories of most of the major distros.



# Your Daily Dose of TECHNOLOGY NEWS



## IF IT'S TECHNOLOGY IT'S HERE

### DAILY TECH UPDATES >>>

**Free  
Subscription**

Simply visit [www.efytimes.com](http://www.efytimes.com)  
and **subscribe** to the newsletter

From the EFY Group on subjects ranging from nanotechnology to open source software can land in your mailbox for **FREE!**

Hardware, Software, Consumer Electronics, Robotics, Automotives, Telecom, Space—any news that's related to technology will be featured here.

# [www.efytimes.com](http://www.efytimes.com)



**EFY Enterprises Pvt Ltd**

D-87/1, Okhla Industrial Area, Phase I, New Delhi 110 020; Ph: +91-11-26810601-03  
Fax: +91-11-26817563, 26812312; E-mail: [info@efyindia.com](mailto:info@efyindia.com); Website: [www.efyindia.com](http://www.efyindia.com)



# New Horizons KNOW HOW



ONA@NHINDIA.COM

senior corporate  
manager—technical  
**Vipin Sharma**  
New Horizons India Ltd

a look at <http://rpm.livna.org/rlowiki/>

From the main page, download the 'Fedora 9 repository RPM'. This will download a package called *livna-release-9.rpm*. Install it as the root user as follows:

```
rpm -Uvh livna-release-9.rpm
```

Now, simply execute the following command:

```
yum -y install mplayer gstreamer-  
plugins-bad \  
gstreamer-plugins-ugly xine-lib-extras-  
nonfree
```

That's it! You can now enjoy a complete multimedia experience on your Fedora system.

**Q I have installed openSUSE 10.3 on my laptop and it works perfectly. I have a few commands that I need to run as the root every time I log in to my laptop. Is there any way for me to auto-run the set of commands without having to enter them every time? Please suggest how to do that, if it's at all possible?**

—*Vandana Sharma, Faridabad*

As you have not mentioned the commands that you need to run on every log-in, I would suggest you create a script and set it to auto-run every time you log in. Here are the steps that will help you to do this. Remember to become the root user before doing so.

Open a terminal, go to the */etc/init.d/* directory and create a file called *myscript*. Now open the file in any text editor and enter the command that you want to run at start-up and save it. Following this,

make it executable as follows:

```
chmod +x myscript
```

And finally:

```
chkconfig -a myscript
```

Now restart your computer and see the script auto-run at start-up. Hope this solves your problem.

**Q I have an old installation of Fedora Core 6 on my system. It has Firefox 1.5 installed. Please let me know if there is any way by which I can create a new profile for Firefox. I have created different profiles for Firefox on my Windows computer using a profile manager. Have tried the same on Linux as well, but was unable to do that. Please help!**

—*Madhur Shivraj, Dehradun*

It is definitely possible to create a new profile for Firefox on Linux too. To do so, you need to close the application and make sure that it is not running even in the background. Now open the terminal and go to the Firefox program directory and then execute the following:

```
./firefox -profilemanager
```

A profile manager similar to the one that you used on Windows, will open. Once you click the *Create Profile* button, you will be provided with a wizard to create a new profile. Type the name of your new profile and hit *Finish*. This will create a new profile for Firefox. Remember to uncheck the option that says, "Don't ask at start-up." Now you can decide which profile to use while starting up your Firefox Web browser.



**Q I've installed Fedora-9 on my system. Seeing that it is much faster and more secure than Windows, a few of my friends requested me to install the same on their system too. After installation, everything worked fine, but my friends and I are unable to play any audio or video files. I have downloaded the required codecs for the player, but do not know how to install them. Please help as I am new to Linux and do not want to switch back to Windows. I would also like to know what, "...swap partition should be twice the size of the main memory," means?**

—*Amit Jha, Ranchi*

Fedora has a policy to keep its distro free of any patent-encumbered software codecs, which typically include MP3 and various other audio/video codes. So, you won't get support for these directly from the Fedora distribution. However, there's a third party software repository that provides support for multimedia codes on Fedora by making the suitable packages available. Have

# New Horizons KNOW HOW



## Whenever you need any top-notch IT certification, go for 'New Horizons'!

New Horizons, one of the world's leading IT training companies worth over USD 350 million, has extended its global alliances to India to partner some of the world's leading technology companies. For the corporates, professionals and students, this means that New Horizons' learning experience is further complemented by its faculty, comprising of highly skilled and qualified trainers. As the pioneer of revolutionary learning methods and with a global network in over 60 countries, New Horizons' **Corporate Training**, **Vendor Certification Programs** and **Career Development Programs** address the needs of corporates, professionals and students to retain and acquire new technology skills that would keep them abreast with the evolving world of IT.

New Horizons has been conferred the 'Best Training Partner, Northern Region' Award by **Redhat** for four consecutive years.

For any of your top-notch IT certification needs, e-mail at [info@nhindia.com](mailto:info@nhindia.com) And, for corporate enquiries e-mail at [corporate@nhindia.com](mailto:corporate@nhindia.com)

### Vendor Certification Programs



### Career Development Programs



\*Conditions apply

New Delhi: NBCC Place, Lodhi Road (adjacent to JLN Stadium). Tel.: 24366316, 24365715, 24366349.

India's fastest growing Training & Education brand

[www.nhindia.com](http://www.nhindia.com)



# Miro

# Amarok for Video

The one-stop shop for all your video needs.

**A**ccording to the answers posted on the Miro site, in response to some frequently asked questions, "Miro is a free application that turns your computer into an Internet TV video player." However, it has the capabilities to be much more than just an 'Internet TV video player'. It provides a one-stop shop for all your video needs. Its top features include the Miro

Guide for channel surfing; watching folders for new videos; full torrent support, so you can download and view torrents in the same app; resumable playback; video sharing and hosting.

First of all, let's understand what they mean by an 'Internet TV video player'. The concept is simple: it's a video player that can subscribe to and download video podcasts while comprehensively managing them. Sounds similar to iTunes, right? But Miro has many other add-ons, including BitTorrent support, to distinguish itself. Moreover, there is no iTunes for Linux (at least, not yet), so Miro is your best bet.

Well, if you have never used iTunes before and vodcasts (short for video podcasts) sounds alien to you, here's a short introduction. To put it simply, podcasts/vodcasts are nothing more than RSS feeds for audio/video content. So, like any RSS feed, you first need to subscribe to it using an application that understands it—for example, iTunes, Miro, Banshee, Amarok, et al. What these apps do is download new content

## MIRO ON THE FORMAT WARS

In response to FAQ, this is what the site has to say: "We strongly believe that format wars among commercial entities have been a huge stumbling block to advancing Internet video. The best way out of the 'format wars' is to support as many formats as possible and users shouldn't have to think about formats at all. We will be adding support on a continuing basis for AVIs, Flash, Real, and Windows Media."

"In terms of open source, patent-unencumbered codecs like Theora, our goal is to support them as soon as we can, and once open source media players and publishing tools get a bit more solid and commonplace, to nudge publishers to use them."

to your hard disk as and when it becomes available, which you can then listen to or watch at your own leisure—either on the computer or on portable media devices like an iPod.

## Social channels

The hard part of subscribing to videos is finding feeds you are interested in and Miro makes this job very easy for you with 'The Miro Guide'. The built-in and Web accessible Miro Guide [[www.miroguide.com](http://www.miroguide.com)] is a full-featured Web service that provides a comprehensive catalogue of video RSS feeds anywhere (Figure 1). With over 5,000 channels and growing, thanks to an active social community, it's pretty safe to say that you will find something to satisfy your requirements, whatever they may be. Miro comes loaded with a couple of starter channels like the Wired Science video podcast, NASA's jet propulsion laboratory, and even a channel that teaches you how to use Miro, called 'Using Miro'!

Like Last.fm and any other Web 2.0 service, Miro also works by using the data generated by users. Miro Guide is an open directory, that is, anyone can submit an RSS feed, thus ensuring more channel options for users. Also, like Last.fm, Netflix or Amazon, users can rate any channel in the Guide and get suggestions about channels they might like. The best thing about the Miro Guide is that it doesn't lock you in—anyone can create an alternative guide of videos and feeds that you can add to Miro.

Apart from the option to set any channel to download new videos as soon as they are published, you can also instruct channels to stop downloading new stuff if unwatched videos are piling up. See Figure 2. You can even start individual video downloads by pasting the URL in the menu item. Plus, you can easily pause and resume any individual download, or all downloads in one channel, or even all Miro downloads!

## Seamless BitTorrent

Miro can download individual BitTorrent files and torrents that are in feeds. When a video torrent is downloaded, it will be in your channel and library, ready for you to watch, just like any other video download.

## Play any video

This is probably the best part of Miro. The formats it supports vary across platforms, but generally it can play almost all the major formats, like MPEG, Quicktime, AVI, H.264, DivX, Windows Media, Flash Video, etc. To quote from the Miro FAQ page, "The Linux version of Miro uses GStreamer or Xine to play videos. Xine supports MPEG 1/2/4, DivX 3/4/5, Windows Media 7/8, QuickTime, Theora, and more... GStreamer has varying levels of support: Theora and AVI are well-supported; MPEG 1 is supported but has licensing issues; AAC and H.264 are not well-supported..."

A nice thing about Miro is that it works so well for HD content, you'll find lots of video that looks beautiful in full screen, even on the largest displays. As they say on the homepage, 'More HD than anyone!' Figure 3 shows Miro playing HD content from DiggNation's HD video



Figure 1: The built in Miro Guide, a full-featured Web service that provides a comprehensive catalogue of video RSS feeds

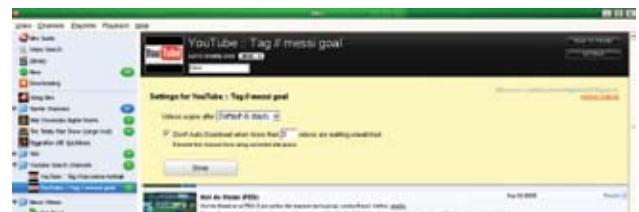


Figure 2: Individual settings per video feed

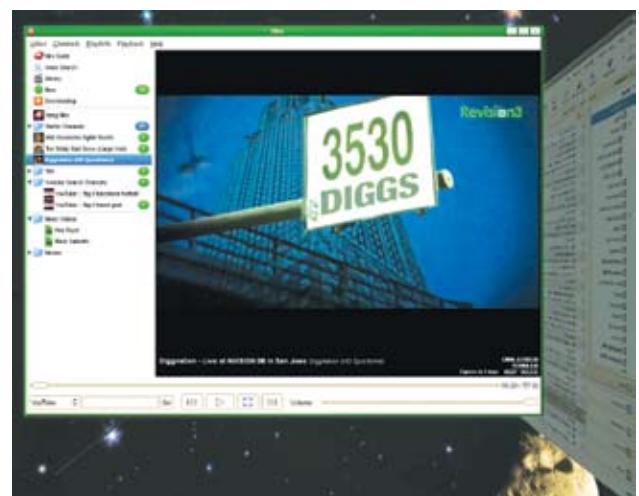


Figure 3: Miro playing DiggNation's HD video

Another really useful feature is the option to 'resume from where you left off'. You can let Miro remember where you stopped watching a video and start at that point when you play it again. Plus you have keyboard shortcuts as well for all key playback commands. Go to *Video*→*Options* in the *Menu* to configure Miro the way you want it.

## Watch a folder

Whenever you add new video files to the 'watched folder', it shows up as a green bubble (look at the left hand side of Figure 1) against the folder channel name. For a movie/music video buff like me, this is very helpful. As with many



Figure 4: Saving a search criteria as a channel

people, I've the habit of collecting, no make it 'hoarding', movies, TV series, music videos, etc, planning to watch them, only to end up forgetting to. But now with Miro, the green bubble is always there as a reminder showing how many videos I've not watched as yet—really helpful to clear out one's video backlog, I should say!

## Organise and manage your videos

Like Amarok, the first time you launch Miro, it asks whether you want it to search for video files in the computer and add them to your library. This is a really useful feature if you ask me—you may end up finding videos you didn't even know you had!

Afterwards, you can create video playlists by dragging them one by one, or selecting a few and right-clicking to add to a playlist. You can even group the playlists and channels into folders for better organisation. For example, you can create playlists like Pink Floyd, The Doors, etc, and group them under a folder called 'Classic Rock Videos'.

Miro also does a good job of HD space management. For example, you can tell Miro to reserve a certain amount of space on your hard drive and it will stop downloading new videos when it reaches that limit. If you are running out of space, Miro can even move your video collection to some other location on your system, such as an external hard drive. Pretty neat, huh?

## Searching

Apart from the channels you have subscribed to from the Miro Guide, you can also search and download from the biggest video sites, including YouTube, Yahoo, Google Video, Blip and more, from right within Miro—and extra brownie points for reducing the number of clicks by including the search box at the bottom in the main window itself.

## Create search channels

You can save any search criteria itself as a channel. Go to *Channels*→*New Search Channel*; in the pop up, search for *field* (see Figure 4). For example, enter 'Messi goal' and select your preferred search engine; say YouTube, and click *Save*. So, as and when someone uploads a Messi goal video

on YouTube, the search channel gets updated. For people who use the Flock Web browser, this may seem familiar—the only difference is that Miro automatically downloads the videos onto your hard disk.

## Search within a channel

You can even save any search within a video feed as its own channel. If you want to auto-download BoingBoing TV, but only when they mention 'steampunks', Miro makes it easy! In the above *New Search Channel* pop up, just select the channel on which you want to search instead of a search engine, enter your search criteria and click 'Create Channel'.

The only gripes I had with searching were:

- The YouTube search is very slow! Sometimes it is better to search in Google, get the URL, and then paste and download it as an individual item.
- YouTube search doesn't show the rating against the videos, so you have no way of knowing which video is better to download.

## Sharing

As I mentioned in the beginning, Miro is pretty much geared towards being a Web 2.0 application. It has wonderful support for sharing. Every channel and video has an e-mail button to quickly send a link to a friend, links to post to Digg, Reddit, del.icio.us, etc. Apart from this, you can even export your feeds as an OPML file and your friends can import it into their Miro or any other feed reader. Miro also provides assistance in creating and publishing video channels.

I think another real nice feature to have, would be IMDB [[imdb.com](http://imdb.com)] and flixster [[flixster.com](http://flixster.com)] integration, as going forward, more and more people will be using Miro to watch and manage their movie collection.

For all the wonderful stuff Miro does, there's still one major shortcoming, which is its inability to sync iPod/iTouch/iPhone, when the main idea behind podcasts/vodcasts is to play the content on these sorts of devices. Among the answers to FAQ, it is said that they will be working on this in the future versions, but when will that become 'present' is what I wonder. In the meantime, Banshee's latest version already supports video, video podcasts, and syncing videos to iPod as well! Maybe it's time for me to give Banshee a spin?

However, all said and done, I'm hoping that as more and more people start using Miro, this requirement will be quickly addressed. With a tagline like "Free and open source, because open media matters," I would really like to see this one succeed.

As for the 'present', the more I use Miro, the more I find it to be the best tool to organise and manage my video collection. Guess it won't surprise anyone if I say it has become my default video player, replacing VLC and MPlayer. 

**By Puthali H.B.** The author is a programmer at Novell, who loves music and open source. To know more about what she finds interesting these days, go to <http://puthali.googlepages.com>

Get more out of IT - Virtual Infrastructure

**Server Consolidation,****Containment,****with VMware Virtual Infrastructure****Business Continuity,****And Disaster Recovery,**

**Experience the difference in learning  
virtualization from the pioneers.**

**G.T. Enterprises has the credit of being the..... (In India) :-**

- ☞ First VMware Partner since 1998.
- ☞ First VMware Authorised Consultants.
- ☞ First VMware Authorised Training Partner.
- ☞ First VMware Premier Partner.
- ☞ First to have received VMware Market Maker Award for APAC for two consecutive years 2005 & 2006.

**The VMware Certified Training Program offers technology professionals the knowledge, skills, and credentials to deploy and maintain virtualization technology.**

**Benefits of VMware Professional Training :-**

- ☞ VMware Infrastructure delivers measurable savings in both capital and operating costs.
- ☞ Demonstrate technical expertise to your customers.
- ☞ Enhance skills to sell, deploy and service VMware products.
- ☞ Ensure superior customer service and support of VMware products.
- ☞ Helps accelerate the application development and deployment lifecycles.
- ☞ Learn from VMware **Gurus** and enjoy the post training extended lab facility.

**Seats are limited!! Hurry**

**We offer courses on VMware Virtualization 3 :-**

- VI 3 : Install & Configure 3.5 - (4 Days) .
- VI 3 : Deploy, Secure & Analyze 3.5 - (4 Days).
- VI 3 : Virtual Desktop Manager - (4 Days).
- VI 3 : Fast Track 3.5 - (5 Days).

**For more details contact us :-**

Email : [training@gte-india.com](mailto:training@gte-india.com)  
Cell Phone : +91-9880223289, +91-9845227221  
Tel : +91-80-26695890-94 (05 Lines), Fax : +91-80-26695887  
URL : [www.gte-india.com](http://www.gte-india.com)

**G.T. Enterprises**

"G.T.House", #48, 1st "B" Cross, 7th Block, Bhavani Layout, B.S.K. 3rd Stage,  
Bangalore - 560 085, INDIA

**TRAINING LOCATIONS :-**

• Bangalore • Chennai • Gurgaon • Hyderabad • Mumbai • Mysore • New Delhi • Pune.

\* Copyright Note : All brand names and product names are registered trademarks of their respective owners.

# FastMail



## For Those Who Hate To Wait!

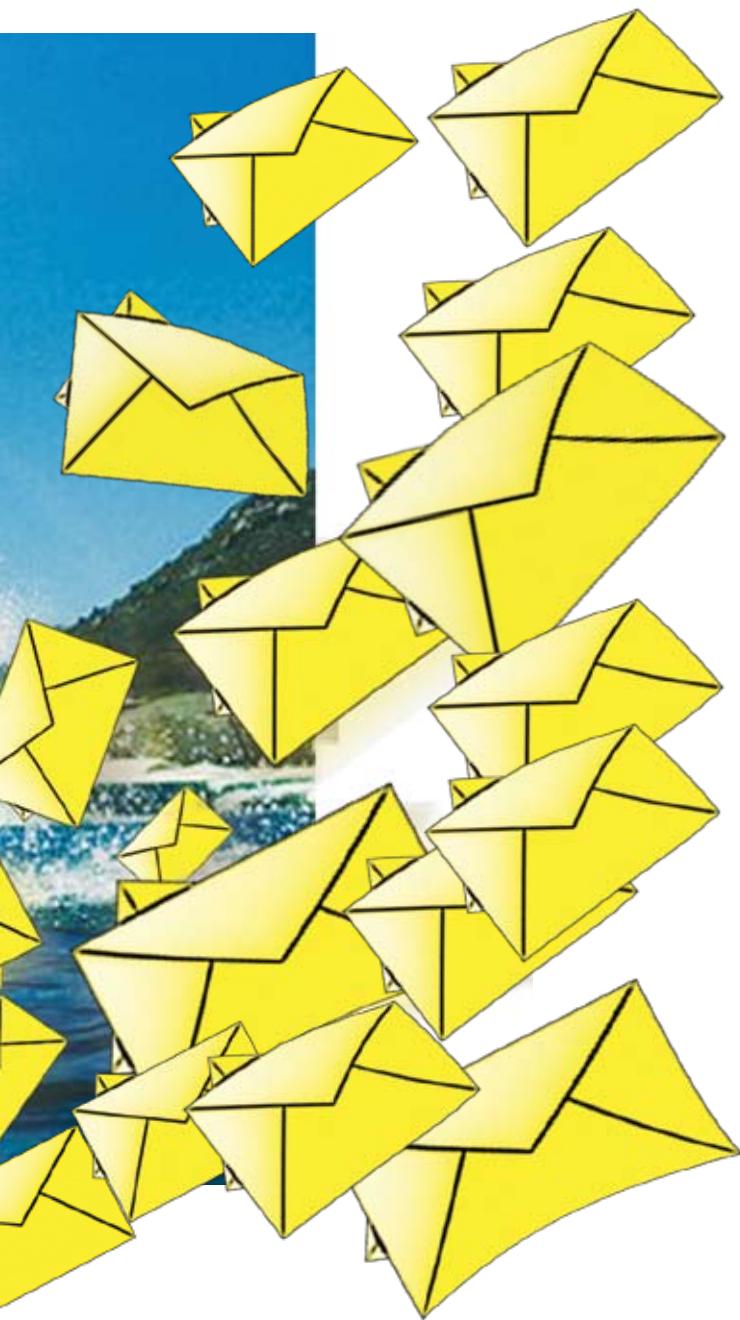
Have you dared to try the power-packed e-mail service that loves to flaunt its FOSS power?

**M**ost people probably haven't heard of the FastMail.fm e-mail service, but those who have will tell you it's named so for a reason. FastMail will already have completed loading your inbox in the time that it takes most better-known e-mail services to ask you to 'please wait' for the service to load. Interestingly, this service runs almost entirely on free software, and seems to have absorbed a certain amount of the FOSS community's spirit. It engages with its users in a bustling, and frank online forum in the style of a GNU/Linux distribution. It's definitely worth taking a look at this intriguing

enterprise. We'll do an overview and test the service.

Established in 1999, FastMail [[www.fastmail.fm](http://www.fastmail.fm)] is a trust from the land of speedsters like Brett Lee, it's Aussie mate, based in Melbourne. The *.fm* root domain is from the Federated Republic of Micronesia, a cluster of islands in the Pacific Ocean, while FastMail's hardware set-up is located in the US.

FastMail promotes itself by offering free accounts, and it's in the business of selling larger and better-featured paid ones to families and enterprises. It will also offer you a choice of domain names. At the time of writing, apart



from [yourname@fastmail.fm](mailto:yourname@fastmail.fm), there's [@rushpost.com](http://rushpost.com), [@123mail.org](http://123mail.org), [@speedpost.net](http://speedpost.net), etc—in fact, there are 104 in all. Of these, most domains are also available for free accounts. FastMail also seems to make a trickle by displaying advertisements in the Web mail interface.

It's a small enterprise. Apart from Rob Mueller and Jeremy Howard, there are two non-executive directors, and a few 'contract programmers', according to the FastMail documentation.

## Openness in software and business

FastMail uses 'nearly 100 per cent' FOSS software, says founder Jeremy Howard. 'I say 'nearly' since IBM only provides binary versions of their RAID configuration and

monitoring tool... but I think that's the only piece of closed source software we use,' he adds. FastMail uses a Debian-based distribution around a custom-patched Linux kernel, a kernel-level firewall, an Apache Web server, Cyrus IMAP/POP server, Postfix 'secure mailer' and a Perl-based main application, besides SpamAssassin and ClamAV.

When asked, Howard does not go into the figures of the business, but shares that use of FOSS has saved the service "...hundreds of thousands of dollars". However, "That's not the main reason we use it," he says. "We use open source software because it is far better for our needs than the alternatives... We often use cutting-edge hardware to get the best possible performance out of our infrastructure. We've worked closely with Linux kernel developers to ensure that it is tuned correctly for this hardware. We'd never be able to do that with Windows!"

Another benefit of FOSS is that the FastMail team can patch its software itself, instead of begging for the same service from a vendor. It has contributed several patches to the Cyrus IMAP server, for instance. That makes its software more reliable, Howard says. "In summary: if we used closed source software, our expenses would be higher, and our uptime would be lower..." It is, he concludes, good for business.

FastMail seems to have—excuse the cliché, but—a philosophy of openness. They've mentioned, in a part of their profuse online documentation, all the details of their hardware, right down to their cabinet map. It's published online at [nyi.brong.fastmail.fm/cabinets.html](http://nyi.brong.fastmail.fm/cabinets.html). FastMail uses IBM x345/x346 and x3550, and Sun's x4500 servers; they have even uploaded photos of their cabinets online at [cabinets.robm.fastmail.fm](http://cabinets.robm.fastmail.fm). Simple gestures like these make FastMail feel more like a FOSS community project than a stuffy e-mail service provider. Should I mention the lively wiki community gathered around it (at [www.emaiIdiscussions.com/forumdisplay.php?forumid=27](http://www.emaiIdiscussions.com/forumdisplay.php?forumid=27)), which uses that space to post queries and suggest improvements directly to the FastMail staff?

## The e-mail service: A review

I tested FastMail's free e-mail service. FastMail lives up to its name; on my Sify Broadband connection of 128 KBps at the most, FastMail usually displays the inbox within seconds, while GMail is still 'Loading'. The reason is that FastMail is devoid of Ajax, Flash, and other technology that's hard on the bandwidth. Its interface is well-endowed and uncluttered, but best described as bare. If you've ever used e-mail in the early days of its popularity in India, you'll get a mental flashback by using FastMail. But FastMail has modern features. Besides e-mail, it offers file storage, a notepad, an address book, photo-and-file upload and other options, all nicely laid out in tabs (Figure 1).

Few e-mail interfaces are as well thought out as FastMail's—perhaps user feedback is to be credited here. Every common mail-related option is laid out on screen. The add-attachment box, subject box, and the CC and BCC

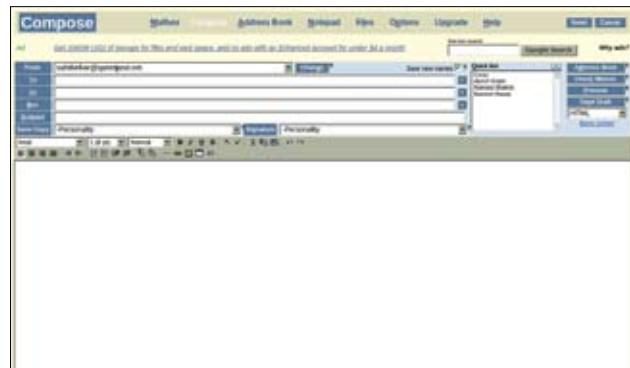


Figure 1: The compose mail interface offered by FastMail

text boxes are open by default. Recently-used addresses are available in a quick-view pane. While most e-mail services only display a few words, FastMail displays the first few 'lines' of unread messages below them, so you needn't click open untitled e-mails to find out what they say. All these points seem trivial until they save you irksome mouse-clicks.

The interface is not pretty, but you can customise it. There are many colour themes to choose from, some of them user-contributed.

Speed is the word at the back-end too; FastMail claims it queues no mail, and that all messages are sent immediately. An excerpt from the FastMail documentation says: "Our pages are generated by our Web application server in 1/100th of a second. They are sent through a 100 MBps link that has plenty of spare capacity. They then go out to the Internet through network links of seven of the most reliable and most fast backbone providers ... we never have any mail queue (all mail is delivered within seconds)."

So far so good, till we came across a sore point. If you use a free account, FastMail embeds a self-promoting text tagline at the bottom of your sent messages. The tagline is text-only -- no flashing GIFs or embedded links. The e-mail interface does show advertisements above your inbox, but so far I haven't come across the typical loud ones, only an unobtrusive line of text or a sponsored link. FastMail doesn't serve you targeted advertisements (like GMail) either.

Having used FastMail a while, you'll suddenly notice the absence of a spam folder in your inbox. Mail from known spammers is blocked automatically at the server level with FastMail's custom filters. That is, for free accounts, spam—or what's detected as spam—isn't delivered to the inbox at all. Paid accounts get the additional benefit of SpamAssassin, an Apache product, for which, FastMail claims a 95 per cent spam blockage rate; this can be customised by the user for greater accuracy. For all accounts, free or paid, images embedded in e-mails are blanked out by default, to throw off spammers, who might have linked to them.

What about security? FastMail leaves no cookies on your computer, and doesn't use Java or Javascripts. In keeping with its FOSS policy, FastMail has implemented the Clam anti-virus, which is free/libre; but there are also a lot of opinions online that it is less effective with viruses than the corporate, paid-for anti-virus software.

In its documentation, though, the mail service defends its use of ClamAV thus: "The best e-mail gateway anti-virus software. Don't fall for the claims of anti-virus software vendors. They specialise in anti-virus software running on Windows machines, not in detecting viruses in e-mails, which ClamAV does better than any of the commercial products." It's an intriguing claim.

And now, the sobering news. So far we've been going rah-rah over FastMail, but, depending upon your e-mail habits, there're a few points that need to be pondered over really hard.

If prolonged GMail use has given you the habit of never deleting messages, no matter how dated or banal, you won't like FastMail a smidgeon. It gives its free accounts—ready?—10 MB space each. (Only paid accounts get 6GB or less, depending on tariff scales at [www.fastmail.fm/pages/fastmail/docs/pricingtbl.html](http://www.fastmail.fm/pages/fastmail/docs/pricingtbl.html).) *What? In this age?* But there are still people who like that constrained e-mail space: it forces them to clear out the garbage.

Moreover, to trip up spammers, FastMail imposes a limit of 80 messages sent per hour, exceeding which it temporarily freezes your (free) account. Also, there's a 40 MB monthly transfer quota for sent messages plus attachments. This is enough for the average free account user, as FastMail demonstrates in its documentation (divide 40 MB with, say, 50 KB, which is the average size of an e-mail). This includes attachments; which must, moreover, not exceed 10 MB, sent or received.

Even receiving certain e-mails is disallowed. Specifically, image-rich e-mails from Yahoo Groups, because they apparently take up too much space.

Finally, FastMail doesn't give free account holders access to its SMTP servers; to use FastMail with an e-mail client like Thunderbird, therefore, free account-holders must use their Internet service provider's SMTP server. Otherwise, they must use the Web interface.

Private users might not mind this hobbling, but companies might. So, this hobbling is intended to get them to buy the upgraded service.

## Final cut

FastMail is the very opposite of GMail or Inbox. It doesn't offer a Texan ranch for space; there's a little space, and you've to make regular deletions to keep it unoccupied. The reason to use the free account, then, is its power and simplicity, and its free/open philosophy. And despite FastMail's hobblings and nobblings, there's a refreshing mom-and-pop corner shop cosiness in its complete transparency. Here is a service, you feel, that tells you everything; and that won't land you in a soup. Not that others will, but you're especially reassured with this one. Besides being fast and feature-rich, it uses free software. Well, I'll be keeping my account! 

**By: Suhit Kelkar** is a freelance journalist based in Mumbai. He uses GNU/Linux.



## EPIA NX-Series Nano-ITX Motherboard

Low Profile & Ultra Compact  
for Flexible System Design

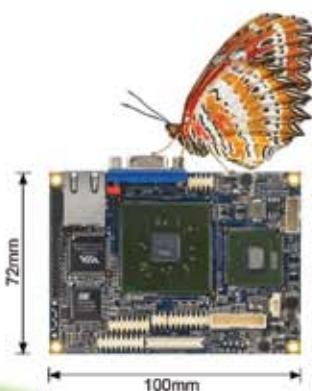


LVDS-08

- Dual view support on 2 dual channel LVDS panels with one additional add-on card (LVDS-08)
- VIA C7 1.5GHz or VIA Eden Fanless 1.2GHz Processor
- DDR2 400/533, SODIMM
- MPEG-2/ MPEG-4/ WMV9 Hardware Decoding Acceleration
- VIA Vinyl HD Audio
- VIA TV Out with HDTV resolution

## EPIA PX-Series Pico-ITX Motherboard

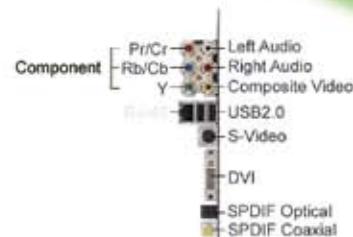
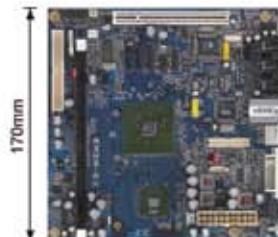
The World's SMALLEST X86 Embedded SBC  
for Networking, Connectivity, Multimedia



- Revolutionary Form Factor
- VIA C7 1.0GHz Processor
- DDR2 400/533, SODIMM
- Pin Header I/O for flexibility
- MPEG-2/ MPEG-4/ WMV9 Decoding
- High Definition Audio
- HDTV Resolution TV-out optional

## EPIA EX-Series Mini-ITX Motherboard

The Ultimate Digital Media Device Platform



- VIA C7 1.5GHz or Fanless 1.2GHz Processor
- DDR2 400/533
- MPEG-2/ MPEG-4/ WMV9 Hardware Decoding Acceleration
- S-Video and Composite TV-out
- Component jacks for HDTV output
- 8-ch S/PDIF output audio experience; 2-ch HD Audio with RCA I/O
- Integrated DVI / LVDS connectors

AHMEDABAD Tel.: +91 - 79 - 2687 1977 / 2988  
BANGALORE Tel.: +91 - 80 - 2527 4390 / 2525 0220  
HYDERABAD Tel.: +91 - 40 - 2774 1122 / 2772  
NEW DELHI Tel.: +91 - 11 - 2671 2358  
PUNE Tel.: +91 - 20 - 2563 1081 / 6601 1818



# Software Freedom Day

## The Delhi Chapter!

An eye witness account of the Software Freedom Day celebrations in Delhi!

**S**oftware Freedom Day celebrates the spirit of freedom in software and introduces it to those who are still trapped inside the non-free software jail. According to the Software Freedom Day website [[www.softwarefreedomday.org](http://www.softwarefreedomday.org)], it dates back to August 28, 2004, when over 70 teams participated in the celebration. However, it was in 2006, when people decided that the Software Freedom Day would be

held on the third Saturday of September, annually—this year it was September 20! September is also the month when Richard M Stallman, a.k.a. RMS, conceptualised the idea of GNU; thus the month also celebrates the birthday of GNU!

### **Fresh rain, free spirits and Delhi**

It was drizzling in the morning and there were apprehensions about heavier rain. Would it wash out the celebrations? However,

all that the little bit of rain did was make the hot summer day a bit cooler—a perfect setting for celebrations.

Folks started gathering at Sarai, the venue of the event. Once most of the key presenters and attendees were in the CSDS (Centre for the Study of Developing Societies) seminar room, Gora Mohanti gave the green signal to start the event. It took off with a very brief introduction about the Software Freedom Day, and then Raj Mathur took over the helm.

## Licence Raj

Mathur took the audience back into the history of free software—how Richard M Stallman stirred the Free Software storm from MIT. He joked about how free software came into existence just because a programmer was too lazy to go to the printer and check if paper had jammed it. He mentioned other important milestones on the road to freedom and touched upon issues related to licensing.

Later, Mathur was queried about his statement that Free Software can be used by anyone, for good as well as bad use, without any restrictions—just like a knife that can be used to cut fruit as well as slit a throat. He was asked that when there are awesome mechanisms like the GNU GPL and other licences to stop abuse and misuse of free software by proprietary companies, why could a similar mechanism to stop misuse by criminals and terrorists not be set up?

The discussion then moved ahead. When Mathur said that FSF promotes OGG, instead of MP3, one of the attendees pointed out that there are FSF approved GNU/Linux distros like BLAG, which come with MP3 support out-of-the-box, which is provided by the free MP3 decoders.

According to my interpretation, Mathur's point was that FSF only promotes anything that is completely free and shuns patents.

Now, patents are something that restrict many distributions from providing out-of-the-box support for MP3 and other patent-covered codecs. However, since software patents are not recognised in India, as long as you use free software MP3 decoders, you're not doing anything 'wrong'. Mathur also highlighted that countries that recognised software patents are in a minority, pointing out that the most powerful country in the world, as the supporter of software patents, was backward in many senses. The 'backward West'? Good point!

## Nostalgic Kishore

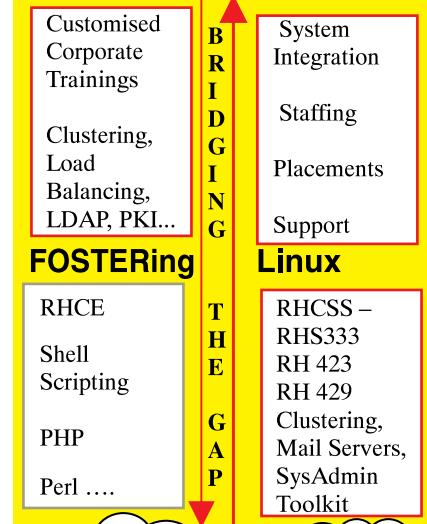
This was followed by a presentation by Kishore Bhargava, on the inception and history of iLUG-Delhi. An interesting fact that he shared was that though iLUG-D members have moved to other cities, in addition to joining local LUGs, they've also started 'regional' chapters of iLUG-Delhi in the cities they moved to. This seems an interesting 'expansion' of iLUG-Delhi across the capital's borders.

Many of Kishore's presentation slides showed LUG members eating out and sharing food, and that surely kindled the appetite of attendees! Luckily, lunch was served right after his presentation.

## LUG@IIT-D and OSScamp

Post-lunch, Gajendra Khanna from IIT-D talked about the work done by LUG@IIT-D till date. He mentioned that the group was formed early this year and had already conducted five workshops. He started a discussion around the need for possible collaboration between members of different LUGs for documentation and projects, especially 'getting-started' kind of documentations.

Kinshuk Sunil of OSSCube used the opportunity to talk about the OSScamps or Uncons that they are going to organise at the end of the month at IIT Delhi. There was some exchange of ideas as to what should be the focus of such events.



**redhat**  
TRAINING PARTNER

Exam Centre for RHCE & RHCSS

**FOSTER** Linux  
Free & Open Source Training Education & Research

SCO- M22, II Floor  
Old DLF Colony  
Sector 14  
Gurgaon  
Tel: 0124-4268187  
4080880,  
9350640169,  
9818478555

55, Deepak Bldg  
13 Nehru Place  
New Delhi - 19  
Tel: 011-30880046,  
30880047,  
9810477448,  
9891074905

Email: [info@fosteringlinux.com](mailto:info@fosteringlinux.com)  
Website: [www.fl.keenable.com](http://www.fl.keenable.com)



There was also a gentleman from Riyadh, Saudi Arabia, who happened to be in Delhi. Dr. Rizwan Ul-haq, a DBA at the King Saud University, was interested in knowing more about the Free Software alternatives available for mainframes. It was interesting to learn that while on the server side they were using SUSE Enterprise Linux, on the desktop front, students in the Saudi university were using various flavours of GNU/Linux.

### Zero hour!

Gora Mohanti slipped back into his role as a speaker with a session that was more like the 'zero hour' at parliament. Various issues were discussed. Topics ranged from the areas LUGs should target, to how to increase the penetration of Free Software among NGOs, educational institutions, and government bodies.

An interesting point made was that small businesses and entrepreneurs could be a big taker of Free Software, as this is a segment which is not only cost sensitive but also less tech savvy—it tends to stick with whatever works for it. So if these businesses are exposed to the benefits of Free Software from the start, they will take forward the legacy.

Increased coverage of Free Software in the mainstream media was the next thing proposed, as mainstream journalists are still not well informed about free software. Then followed a lengthy discussion over whether LUGs ought to start out on some paid-for projects or offer support. There were different opinions on this issue. On one hand, supporters said that there was no harm if there could be a sub-division within iLUG-Delhi that offered paid support for services; on the other hand, others came out with some real-life obstacles in doing that. The topic was reserved for discussion in up-coming events.

### Niyam: Heard loud and clear

Then came the time for Niyam Bhushan to make his presentation on multimedia. But, unfortunately, Audacity could not work on the newly installed Ubuntu Studio. Un-moved and 'gut-sy' Niyam moved ahead with sharing his thoughts on how to 'normalise' the volume of your MP3 collection using Audacity. His tips and tricks included topics like how to save yourself from ear-damage; equalisation presets in Audacity; noise-removal using Audacity, and much more. He also shared an interesting tool called Gnaural with us.

Then the time came for the free distribution of GNU/Linux distros. There were many takers for the Ubuntu 8.04 Desktop editions, while some picked up Fedora 9, to the accompaniment of playful exchange

## LiteratureIndia.com launches Hindi Section

Web magazine 'Literature India' launched its Hindi section [[www.literatureindia.com/hindi](http://www.literatureindia.com/hindi)] on the occasion of the Software Freedom Day and Hindi Fortnight celebrations. The release of the site has become the first programme of its kind, where a Hindi site was launched online through the #sarai IRC channel on Freenode.

The portal was released by noted poet, journalist and translator, Neelabh. The function started with a keynote address by a distinguished historian and writer Ravikant. Ravi Ratalami, a satirist and technical translator, conducted the programme, which was attended by an online gathering of noted personalities.

In his inaugural address, Neelabh identified the relationship between technology and language, and said, "While technical experts are unaware of linguistic knowledge, the Hindi community is afraid of technology."

Ravikant said that bilingualism is itself an important aspect. A Hindi-speaking person generally doesn't write much in English, so this site can work as a bridge. He pointed out that if we search for anything about Hindi writers, we find little about them on the Net. He hoped that this portal would fulfil the aspirations of readers.

In her welcome address, the editor of the portal, Sangeeta Kumari, said that she had been contemplating this portal for the last four years, and only recently the idea took shape. She also added that this site would concentrate on the whole cultural field rather than only on literature.

At the end of this programme, Purnima of Abhivyakti-Anubhuti and Shailesh recited their poems. The programme on IRC lasted for about an hour and 45 minutes. Lots of people from different fields gathered at Sarai's IRC channel and participated in the active discussion.

The website launch can be seen as a success of language technology in open source. This site is based on the Joomla CMS and the server is hosted on GNU/Linux.

of words between the supporters of the two distros.

### All's well that ends well

Finally, dusk was upon us, and it was time to break from what seemed to be a perfect Software Freedom Day. As Gajendra said, "The event successfully met its goal of bringing together various experts to discuss the common issues all of us face on a day to day basis."

Well, what more could a Free Software lover ask for? 



## ATTENTION MCA, MSC (IT) Final Year Students

**GET Trained Project Placed** 

[www.lynusacademy.com/projects](http://www.lynusacademy.com/projects)

- Projects available in LAMP, Java, J2EE, Perl, Python & ROR
- Batch Discount available
- Accommodation will assist For Out Station Students
- Register Now .. Limited Seats Available
- IEEE Projects also available



Authorized  
Training Partner



**Special Offer for  
Job hunters & Freshers  
Zend PHP Training  
( 3 Months)  
Rs. 23,200 / - Only**

### Other courses

- RHCT & RHCE
- Shell Scripting
- Perl Scripting (for Sys. Admin)
- G-FOSS - Web Development (3 Months & 6 Months)
- G-FOSS - Engineer (3 Months & 6 Months)
- Perl + CGI Programming
- Python Programming
- Ruby on Rails
- AJAX & JQuery
- MySQL - for Developer
- MySQL - For Administrator
- PostgreSQL - for Developer
- PostgreSQL - for Administrator



Future is open

**Lynus Academy Pvt Ltd**

Chennai - 044-25021630,  
42171278, 9840880558  
Pondicherry - 0413-2276233,  
4202121, 9952966528  
contactus@lynusacademy.com  
[www.lynusacademy.com](http://www.lynusacademy.com)

### ACKNOWLEDGEMENT

*Kishore Bhargava got trigger happy and shot some awesome photographs on the 20th. All the snaps in this article, except where mentioned otherwise, have been taken by him. Thanks a lot, Kishore :-)*

**By: Swapnil Bhartiya, assistant editor, EFYTimes.com**

*The event was sponsored by LINUX For You magazine and organised in association with Sarai and iLUG-Delhi.*



# A Hitchhiker's Guide to Hypervisors



This article provides a brief overview of virtualisation, with a special focus on hypervisors.



**W**hat is virtualisation? We are familiar with the concept of multiple processes running simultaneously and sharing the resources of a single computer. This is achieved by the operating system acting as a single point of contact interfacing with the hardware resources, and thereby controlling the access to the hardware resources by multiple processes. Virtualisation can be thought of as an extension of this concept, wherein multiple operating systems are allowed to share the hardware simultaneously by means of the virtualisation software.

An operating system directly interacts with hardware in the normal non-virtualised

case. Now if we want to run multiple operating systems simultaneously, we need to have an abstraction layer between the OS and the actual hardware. This hardware abstraction layer fools the operating system into thinking that it is directly interacting with the hardware. The term virtual machine or hardware virtual machine refers to the hardware abstraction layer provided by the virtualisation software, which allows each operating system to think that it is directly accessing the hardware. The virtualisation software that provides this illusion is typically referred to as 'hypervisor' or virtual machine monitor. The terms VMM and hypervisor are typically used interchangeably in this context.

## Types of virtualisation

The two major types of virtualisation techniques are emulation (also known as full virtualisation) and paravirtualisation.

- **Full virtualisation:** Here the virtualisation software provides a complete emulation of the underlying hardware. All software that can run on the underlying hardware can run as is, on the virtual machine. The operating system does not need any modifications to be run as a guest OS instance. The guest OS instance can be any operating system supported by the underlying hardware. VMware workstations, Virtual PC and QEMU are examples of this technique.
- **Para virtualisation:** Here the virtualisation software provides an abstraction, which is very similar, but not completely identical to the underlying hardware. Xen virtual machine and VMware ESX server are examples of this technique. Instead of completely emulating the underlying hardware architecture by the virtualisation software, the virtualised guests collaborate with the hypervisor to achieve optimal performance. Paravirtualisation offers significantly improved performance; however, it requires modification to the guest operating system. The guest OS is modified at load time to include paravirtualisation extensions. Hence, it requires cooperation from the OS vendor.

## Hypervisor basics

Hypervisor is the basic building block for virtualisation. It is an entity that abstracts hardware resources and thus enables running a variety of operating system images concurrently. It's typically built as a derivative of an operating system, as part of platform firmware, or as a standalone embedded solution. It is also referred to interchangeably with Virtual Machine Monitors (VMM).

Hypervisors allow the unmodified, or modified version of the operating systems to be run as guests; sometimes guests intended for an entirely different processor architecture. The hardware resources (like CPU, I/O devices, etc) are virtualised and presented to the guest. The guests run mainly like a user application on the hypervisor. The guest can also work in a cooperative fashion by having modified drivers to avoid performance overheads.

Hypervisors can be of two types, namely, Type 1 and Type 2. In the case of Type 1, the hypervisor runs directly on hardware. ESX server from VMware and Xen are Type 1 hypervisors that sit underneath the operating system on the server hardware. In the case of Type 2, the hypervisor runs on a host OS, which in turn runs directly on the hardware. VMware workstation is an example of the Type 2 hypervisor. Hybrid techniques are also possible, wherein both the host OS and hypervisor can access the underlying hardware. In order to further explain the hypervisor internals, we will focus on the open source Xen hypervisor.

In Xen's hypervisor implementation, the CPU, memory and low-level hardware interrupts are virtualised by a low-level efficient hypervisor layer. When the OS makes

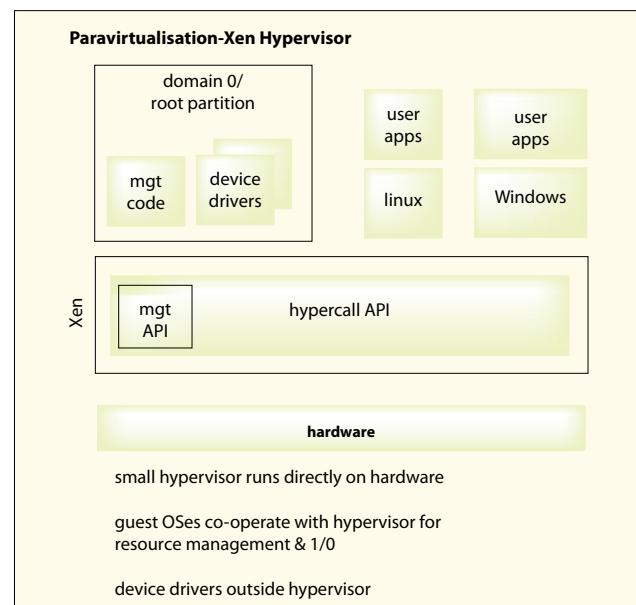


Figure 1: Xen's 'hypercall' interface

changes to hardware-aware data structures, such as the page table, or initiates a DMA operation, it collaborates with the hypervisor by making calls into an API that is offered by the hypervisor. The communication between the guest OS instance, often referred to as the domain in Xen terminology, is by means of this 'hypercall' interface, as shown in Figure 1. 'Hypercall' can be considered similar to the 'system call' mechanism for user processes to tap into the operating system in non-virtualised environments.

The hypervisor is mapped into the address space of each guest operating system. Hence, there is no context-switch overhead between the operating system and the hypervisor on executing a hypercall. Xen makes a guest operating system (running on top of the VMM) virtualisation-aware and presents it with slightly modified x86 architecture, provided through the so-called hypercall API. This removes any difficult and costly-to-emulate privileged instructions and provides equivalent, although not identical, functionality with explicit calls into the hypervisor. The changes needed in the guest OS code to support paravirtualisation are, in general, confined to a few hardware-aware modules, and the bulk of the operating system code and the entirety of application program code remain unmodified.

The Xen hypervisor itself provides only basic control operations. Complex policy decisions for sharing resources between different domains are actually performed by a management software running over a guest OS rather than in hypervisor code. A domain is created at boot time, which is permitted to use the control interface. This initial domain, termed Domain0, is responsible for hosting the application-level management software. The control interface provides the ability to create and terminate other domains and to control their associated scheduling parameters, physical memory allocations, the accessible physical disks and network devices.

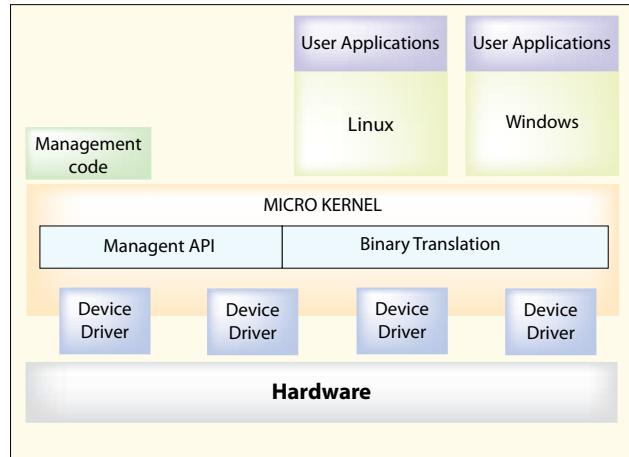


Figure 2: Microkernel virtualisation

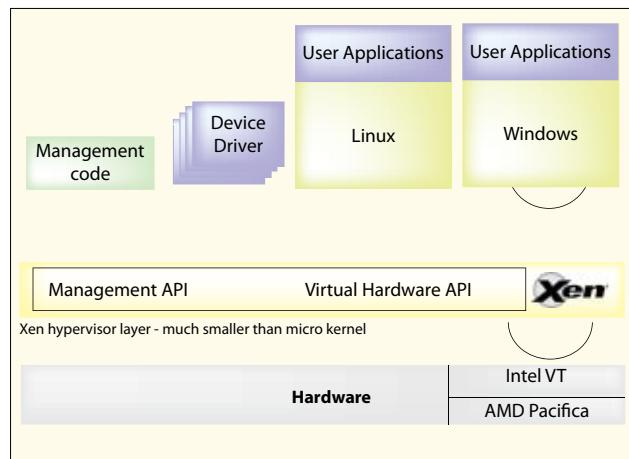


Figure 3: Xen Hypervisor virtualisation

## Virtualisation hypervisor vendors

Citrix (XenSource), Open source Xen, Microsoft (Virtual Server, Virtual PC, Hyper-V), Virtual Iron and VMware are a few major players in the x86 market. HP (Integrity Virtual Machines), IBM (PowerVM, zVM) and Sun (xVM) have proprietary hypervisors on the UNIX market.

**XenServer** system is structured with the hypervisor, and using Dom0 (first guest) the guests (DomU) are hosted. Dom0 can be modified versions of Linux, NetBSD and Solaris. XenServer uses the Xen hypervisor to virtualise each server. It can combine multiple Xen-enabled servers into a resource pool leveraging resource clustering technology. XenServer extends the basic single-server notion of virtualisation to enable seamless virtualisation of multiple servers as a resource pool, whose storage, memory, CPU and networking resources can be dynamically

controlled to deliver optimal performance.

**Xen** hypervisor paravirtualises the hardware, unlike microkernel virtualisation (Figure 2).

This provides much lower performance overhead, especially with I/O. Paravirtualisation alleviates binary patching. The native Linux drivers are leveraged to provide support for a diversity of drivers with a tiny hypervisor code base. Xen's paravirtualised

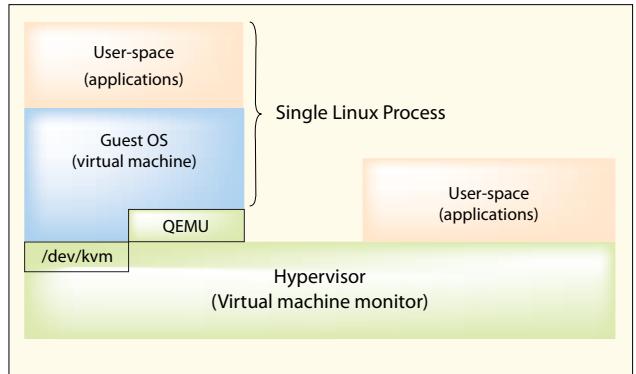


Figure 4: KVM hypervisor

drivers run outside the core hypervisor at a lower protection level than Xen, making the hypervisor impervious to driver failure (Figure 3).

**Kernel Virtual Machine** or KVM is a new Linux subsystem that leverages these virtualisation extensions to add a virtual machine monitor (or hypervisor) capability to Linux. KVM is open source hypervisor software (Figure 4) that provides both full and paravirtualisation capabilities for Linux on x86 hardware containing virtualisation extensions from Intel and AMD. The architecture is optimised to utilise the native functionality of the underlying Linux OS.

KVM is the core hypervisor virtualisation technology that is used in Qumranet's Solid ICE desktop virtualisation solution. KVM uses QEMU (modified), a generic and open source machine emulator and virtualiser. QEMU comprises a *Dynamic Translator* that performs a run time conversion of the target CPU instructions to the host instruction set.

**Sun xVM** (Figure 5) is a bare-metal hypervisor based on the open source Xen under a Solaris environment on x86-64 systems. On SPARC systems, xVM is based on Sun's Logical Domains and Solaris. Sun plans to support Microsoft Windows (on x86-64 systems only), Linux, and Solaris as guest operating systems.

**LinuxOnLinux** is a User-Mode Linux that allows the Linux kernel to run in user space using Linux as a hypervisor for a paravirtualised Linux. One process per virtual processor, some host assist patches, and reuse device approaches are used targeting IA-64.

**VirtualBox** is an x86 virtualisation software package, originally created by German software company Innotek, now developed by Sun Microsystems as part of its Sun xVM virtualisation platform.

**Virtual Iron Software** fully supports Intel-VT and AMD-V hardware-assisted virtualisation. The platform is based on the open source Xen hypervisor.

**HP Integrity Virtual Machine** is a HP-UX based VMM implementation that supports Linux guests in tandem with Windows and HP-UX guests on IA64.

**IBM** provides Linux support on z/VM and PowerVM.

**Oracle VM** is based on the open source hypervisor technology, and supports both Windows and Linux guests.

**Hitachi Virtage** is a hypervisor-type embedded

virtualisation at the hardware layer. It supports Red Hat Enterprise Linux 4 and SUSE Linux 10 on IA architecture.

**Microsoft Hyper-V** has to have at least one parent partition, running Windows Server 2008. It supports SUSE Linux Enterprise Server 10 SP1/SP2. Ubuntu Linux 6.06/6.10/7.10, or Fedora 8/9 is unsupported; however, they have been reported to run.

In the embedded world, the deployment of multiple operating systems on multi-core processor platforms seems the logical step. The **LynxSecure** separation kernel is a bare-metal native hypervisor (Figure 6) intended for use in embedded systems and high assurance security applications for x86 virtualisation.

Military and avionics industries benefit where secure and insecure guests could be mixed. 'Cooperative virtualisation' provides superior performance for the guest operating systems—such as Linux, LynxOS-SE and LynxOS-178.

**RTS Hypervisor** is a hypervisor (Figure 7) said to enable multi-core processors to simultaneously run an RTOS (real-time operating system) and a GPOS (general-purpose operating system)

Individual processor cores, memory and devices can be assigned to guests. In order to facilitate communication between operating systems, the RTS solution also provides a configurable user-shared memory as well as a TCP/IP-based virtual network driver.

**oVirt** is a small host image that provides *libvirt* for virtual machine management, storage management, and secure remote communication. oVirt's pre-built images use the KVM technology built into the main Linux kernel. It would be interesting to observe the development on this front as this could be used as an alternative for Xen or VMware.

## Improving hypervisor performance with hardware assists

Hardware assists come from the processor manufacturers to alleviate hypervisor performance bottlenecks in the process of CPU, I/O, and OS/Platform resource virtualisation. Both Intel and AMD have announced hardware extensions that can assist virtualisation. Hardware support avoids the complications associated with interpretation or binary translation.

Intel's technology is known as VT, which stands for Virtualisation Technology, earlier codenamed as VanderPool. VT-x defines the extensions to the IA-32 Intel architecture, whereas VT-i defines the extensions to the Intel IA-64 for virtualisation support. These extensions are known as Virtual Machine Extensions (VMX). VT-x augments IA-32 with two modes of CPU operations, namely, VMX root operations and VMX non-root operations. The transition from the VMX root operation to VMX non-root operation is known as 'VMEntry'. This refers to the transition from VMM to a guest operating system. The transition from a VMX non-root operation to VMX root operation is known as 'VMExit' and refers to the transition from the guest OS to VMM. By providing hardware support for these costly transitions, the VT extensions help to improve hypervisor performance.

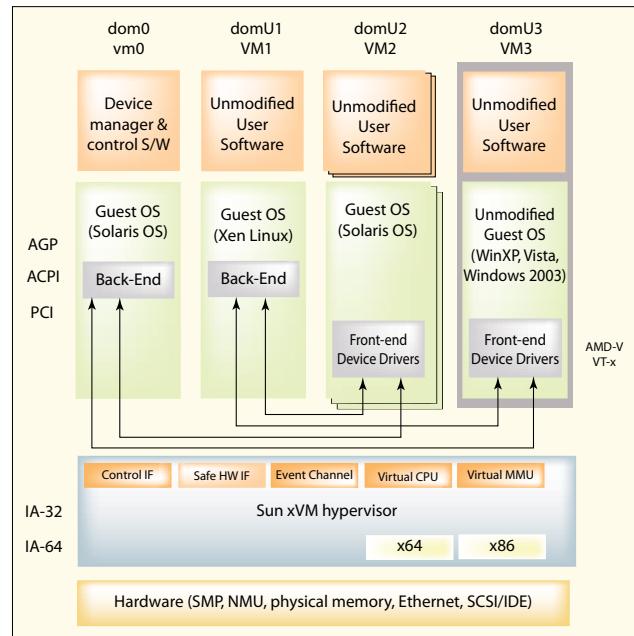


Figure 5: Sun xVM hypervisor

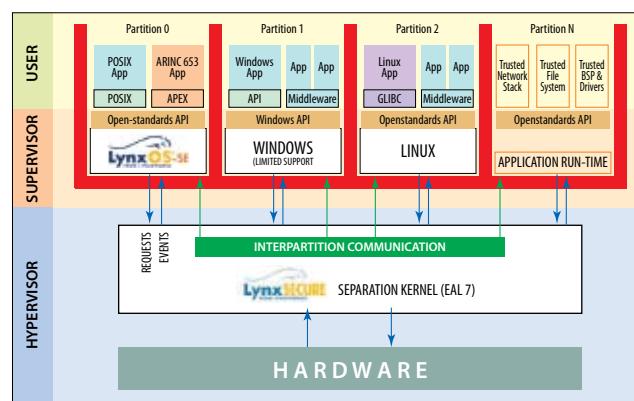


Figure 6: LynxSecure hypervisor

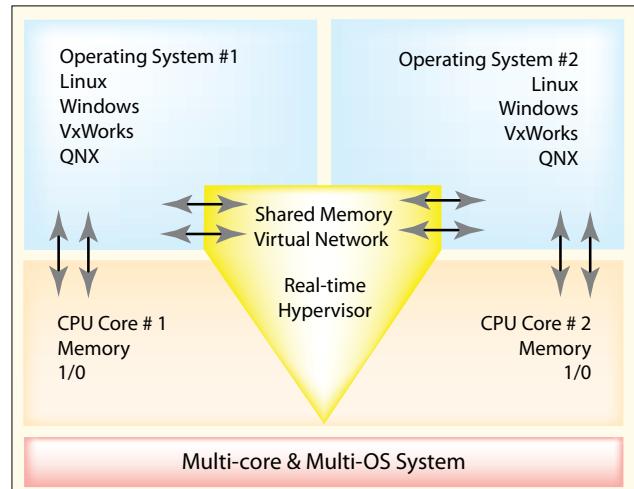


Figure 7: RTS Hypervisor

Intel has also announced further hardware assist extensions to support I/O virtualisation. Intel VT-d hardware assistance provides remapping capability that can be used

**TABLE 1: LIST OF HARDWARE-ASSISTED HYPERVISORS**

Hypervisor	Hardware Support
LynxSecure	VT-x and VT-d
Parallels Workstation	VT-x and AMD-V
Parallels Desktop	VT-x and AMD-V
Parallels Server (Beta)	VT-d
Padded Cell	VT-x and VT-d
VirtualBox	VT-x and AMD-V
Virtual Iron	VT-x and AMD-V
VirtualLogix	VT-x and VT-d
VMware Workstation	VT-x
VMware Server	VT-x and AMD-V
Xen	VT-x and AMD-V (Xen 3.0.2), VT-d (Xen 3.2.1), VT-c (Xen 3.3)

for controlling and monitoring DMA accesses, and also performing direct I/O assignment under the control of the system software. Intel's Virtualisation Technology for Connectivity (VT-c) is itself a collection of technologies that assists in I/O virtualisation. It is composed of Virtual Machine Device Queues (VMDq), Intel I/O Acceleration Technology (I/OAT), and Single Root I/O Virtualisation. More details can be found at [www.intel.com/technology/platform-technology/virtualization/index.htm](http://www.intel.com/technology/platform-technology/virtualization/index.htm).

AMD's virtualisation extensions to the 64-bit x86 architecture are named AMD Virtualisation, abbreviated AMD-V and codenamed Pacifica. AMD-V is present in AMD Athlon 64 and Athlon 64 X2 with family 'F' or 'G' on socket AM2 not 939, Turion 64 X2, Opteron 2nd generation and 3rd generation, Phenom, and all newer processors. AMD has published a specification for a technology named IO Memory Management Unit (IOMMU) to AMD-V. This provides a way of configuring interrupt delivery to individual virtual machines and an IO memory translation unit for preventing a virtual machine from using DMA to break isolation. More details can be found at [www.amd.com/us-en/assets/content\\_type/white\\_papers\\_and\\_tech\\_docs/34434.pdf](http://www.amd.com/us-en/assets/content_type/white_papers_and_tech_docs/34434.pdf)

Table 1 shows a list of hardware-assisted hypervisors.

## Benchmarking hypervisors

Since there were no standard hypervisor benchmarking suites, most hypervisor vendors use a suite of enterprise workloads to benchmark hypervisors. Both XenSource and VMWare have published benchmark results of their hypervisors using a suite of enterprise workloads such as SPECjbb2005, NetPerf, Passmark and a host of other applications. A detailed comparison of the XenEnterprise 3.2 product with VMWare ESX 3.01 hypervisors can be found at [www.xensource.com/Documents/hypervisor\\_performance\\_comparison\\_1\\_0\\_5\\_with\\_esx-data.pdf](http://www.xensource.com/Documents/hypervisor_performance_comparison_1_0_5_with_esx-data.pdf)

There has been considerable effort from various

virtualisation vendors to develop benchmarks for virtualisation. VMMark is a popular virtualisation benchmark suite from VMware, which measures the performance of virtualised servers while running under load on physical hardware. In order to measure the efficiency of the hypervisor, the suite runs several virtual machines simultaneously. Each VM is configured according to a template. The templates mimic typical software applications found in corporate data centres, such as e-mail servers, database servers, and Web servers. The VMmark software collects performance statistics that are relevant to each type of application. When benchmarking, VMs are grouped into logical units called 'tiles'. When evaluating a system's performance, the VMmark software first calculates a score for each tile, culled from the performance statistics produced by each VM, and then aggregates the per-tile scores into a final number.

Intel and IBM have also developed a virtualisation benchmarking suite, namely, VConsolidate. It runs multiple instances of a consolidated database, mail, Web and Java workloads in multiple virtual CPU partitions to simulate real-world server performance in a typical environment. More details on VConsolidate can be found at [www.intel.com/technology/itj/2006/v10i3/7-benchmarking/6-vconsolidate.htm](http://www.intel.com/technology/itj/2006/v10i3/7-benchmarking/6-vconsolidate.htm)

In the embedded industry also, there have been efforts to develop hypervisor benchmarks, for virtualised embedded systems. Embedded Microprocessor Benchmark Consortium (EEMBC) is developing a hypervisor benchmark known as 'hyperbench', which will measure the contribution of hypervisors, to performance, code size, and energy consumption in a wide range of embedded systems. More details can be found at [www.eembc.org/benchmark/hyper\\_sl.php](http://www.eembc.org/benchmark/hyper_sl.php)

## Virtual appliances

Once an operating system and its applications have been encapsulated into a virtual machine, the VM can be run on any computer with a hypervisor. The ability to encapsulate all states, including application and operating-system configuration, into a single, portable, instantly-runnable package provides great flexibility. This is being taken advantage of by means of virtual appliances, which package an application into a virtual machine and can be run anywhere on a computer with a hypervisor. Since virtual appliances are preconfigured, they greatly eliminate the installation, configuration and maintenance costs associated with installing and running complex stacks of software and are widely being used as a software distribution mechanism. More details on virtual appliances can be found at [www.vmware.com/appliances](http://www.vmware.com/appliances)

With more and more vendors joining the virtualisation bandwagon, one interesting question is whether hypervisors will become commoditised. With virtualisation being supported extensively by hardware, that day may not be far off! 

**By: Saravanan Chidambaram.** The author is a specialist in virtualisation and leads the virtual partitions team at Hewlett Packard India. His areas of interest include virtualisation, systems management and cloud computing.



# Virtual Appliance It's 'Just Enough'

If you believe in the 'just enough' factor, maybe virtual appliances are what you're looking for.

**T**o define virtual appliance, I would like to introduce you to the slightly broader world of software appliances. According to Wikipedia, "A software appliance is a software application combined with just enough operating system (JeOS) for it to run optimally on industry standard hardware (typically a server) or in a virtual machine. Software appliances simplify

server applications by minimising the tasks typically associated with installation, configuration and maintenance."

A software appliance is customised to decrease deployment and maintenance time in particular scenarios.

We could now define virtual appliances as software appliances designed to be deployed using a virtualisation solution such as Sun xVM VirtualBox, VMWare, Qemu, etc.

## Virtual appliances vs virtual machines

A virtual appliance is a fully pre-installed, pre-configured application with an operating system environment and is ready for production deployment, whereas a virtual machine is, by itself, without application software. A virtual machine is created with a virtualisation solution and generally gives a sandboxed development environment with the same or different guest operating system than the host operating system.

As opposed to virtual appliances, which are generally domain specific, virtual machines are general-purpose virtual environments set up using virtualisation solutions.

## Giving it a spin

Trying out our virtual appliances is as easy as downloading the appliances and registering them in your favourite virtualisation solution. There are several places to look for virtual appliances depending on the virtualiser you are using. Virtual appliances for VMWare are available at [www.thoughtpolice.co.uk/vmware](http://www.thoughtpolice.co.uk/vmware), while [sourceforge.net/projects/virtualboximage](http://sourceforge.net/projects/virtualboximage) is a place to look for Sun xVM VirtualBox images. However, it is worth mentioning that you can only use the VirtualBox images on the same host on which the image has been built. Also, it involves getting 'dirty' with the VirtualBox command-line tools to register the new image. You will have to use the *VboxManage registervm* command for this. Please read the user manual for VirtualBox to see how to do this.

## Creating a virtual appliance

There are essentially two points that you should keep in mind when you roll out your own appliances:

- Create a virtual machine using your favourite virtualiser and install an operating system into the virtual disk. While choosing the operating system, due care should be taken to include only those bits that are essential to run the appliance. For example, if you want your appliance to function as a Web server, there is no need to have a graphical interface.
- Customise the installation to suit your requirements and distribute the virtual disk and the configuration files.

Ubuntu JeOS (pronounced 'juice') [[www.ubuntu.com/products/whatisubuntu/serveredition/jeos](http://www.ubuntu.com/products/whatisubuntu/serveredition/jeos)] is an efficient variant of the Ubuntu Server operating system, configured specifically for virtual appliances.

You can find out how to use Ubuntu JeOS to develop your own Linux-based virtual appliance if you are a VMware user at [www.linux-mag.com/id/4829](http://www.linux-mag.com/id/4829).

In this article I'll use Sun's VirtualBox to create a virtual appliance using Ubuntu JeOS. This appliance will have the Apache Web server and MySQL server, which will enable you to deploy PHP websites.

Download the Ubuntu JeOS 8.04 image and install it to a virtual disk using Sun xVM VirtualBox. Note that VirtualBox, by default, uses NAT for networking with the guest OS. However, for a typical appliance, we need

a network bridge instead. This is where the game with VirtualBox becomes a bit tricky.

To quote from the VirtualBox user manual: "With Host Interface Networking, VirtualBox creates a new networking interface in software on the host computer. This new software interface will then exist in parallel to your regular interfaces (e.g., on a Linux host, *vbox0* will exist alongside *eth0*). When a guest is using such a new software interface, it looks to the host system as though the guest were physically connected to the interface using a network cable: the host can send data to the guest through that interface and receive data from it. This means that you can set up routing or bridging between the guest and the rest of your network."

The VirtualBox user manual has detailed instructions for setting up host interface networking on different Linux distros, Solaris, as well as Windows. Here I will reproduce the steps for Ubuntu Linux 8.04.

First, install the *bridge-utils* package as follows:

```
sudo apt-get install bridge-utils
```

Edit your */etc/network/interfaces* file to add the following lines:

```
auto br0
iface br0 inet static
  address 10.10.3.4
  bridge_ports eth0
```

Here 'br0' is the name of the network bridge that we are creating.

Now restart the networking services as follows:

```
sudo /etc/init.d/networking restart
```

Next, create a software host interface as follows:

```
sudo VBoxAddIF vbox0 amit br0
```

...where 'amit' is the user who will be running VirtualBox.

Now, modify the network settings of your virtual machine to add the 'vbox0' interface as follows:

```
VBoxManage modifyvm JeOS -hostifdev1 vbox0
```

...where 'JeOS' is your VM's name.

Finally, boot your VM and assign a static IP address to the interface.

Now, with the basic infrastructure in place we shall go ahead to create the virtual appliance. To begin with, update the repository using:

```
sudo apt-get update
```

Install MySQL Server:

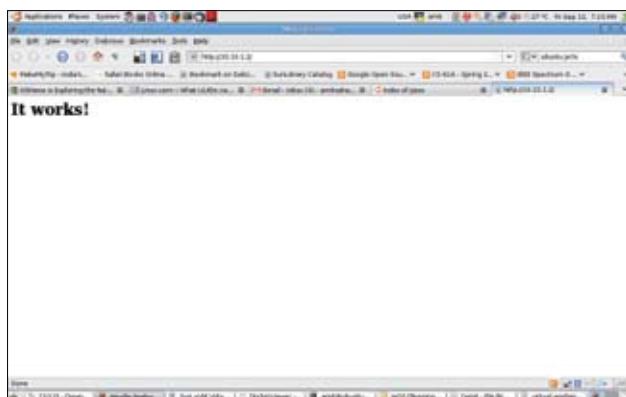


Figure 1: Accessing the Web server virtual appliance from a browser installed in the host machine

```
sudo apt-get install mysql-server
```

Install Apache Web server:

```
sudo apt-get install apache2
```

Install PHP5:

```
sudo apt-get install php5
```

Install PHP5 modules for Apache:

```
sudo apt-get install libapache2-mod-php5
```

Restart Apache:

```
sudo /etc/init.d/apache2 restart
```

You can test it from your host machine. Let us assume the guest IP address is 10.10.1.2. Go to your browser and type: <http://10.10.1.2>. You should see a Web page similar to the one shown in Figure 1.

Note that you can either put your PHP scripts in `/var/www`, or use the MySQL server instead.

That's it! We have got a virtual machine that uses a Linux base and has Apache, PHP and MySQL set up. This is a perfect environment for basic LAMP development. And you guessed it right! We can distribute this as a simple LAMP virtual appliance.

## Distributing your virtual appliance

All virtualisation products use their own formats for the virtual appliances, due to which virtual appliances created using a particular product can only be used properly with that specific solution. This is not a happy situation.

## Open standards to the rescue

A new development that promises to overcome this shortcoming is a new standard for packaging virtual machines called Open Virtualisation Format (OVF) [[www.vmware.com/appliances/learn.ovf.html](http://www.vmware.com/appliances/learn.ovf.html)], conceived by the

Distributed Management Task Force (DMTF). OVF, among other things, will allow interoperability between the various virtualisation products available.

Compared to VMDK or VDI, which encloses only a single virtual disk in the virtual machine, the OVF format provides a complete specification of the virtual machine. This includes the full list of required virtual disks plus the required virtual hardware configuration, including CPU, memory, networking and storage. In short, the OVF is a standards-based portable format that allows the users to deploy their virtual machine in any hypervisor that supports OVF.

OVF makes heavy use of XML and the technical specifications are available at [www.vmware.com/pdf/ovf\\_spec\\_draft.pdf](http://www.vmware.com/pdf/ovf_spec_draft.pdf). More information on OVF is available at [www.vmware.com/appliances/learn/ovf.html](http://www.vmware.com/appliances/learn/ovf.html)

## ovftool

The ovftool is a Java-based experimental tool to convert VMs to and from OVF, and converting standard sparse/flat VMDK files to and from the compressed stream-optimised VMDK format used in OVF. (VMDK is the file format used by VMWare for virtual disks.)

It is available for download at [www.vmware.com/download/eula/ovf\\_eula.html](http://www.vmware.com/download/eula/ovf_eula.html).

## Why care?

According to a Forrester Research report titled, 'The Case for Virtual Appliances': "Virtual appliances enable dramatically simpler on-premise deployments without the burden of hardware management or the infrastructure implications of SaaS. Other benefits include assurance around performance and reliability, simpler administration, lower software life-cycle management, and streamlined upgrades. ISVs should begin evaluating this new option to determine if it can speed deployments and transform their business models for the better." 

## REFERENCES:

- [http://en.wikipedia.org/wiki/Software\\_Appliance](http://en.wikipedia.org/wiki/Software_Appliance)
- [http://en.wikipedia.org/wiki/Virtual\\_machine](http://en.wikipedia.org/wiki/Virtual_machine)
- <http://www.virtualbox.org>
- <http://www.vmware.com>
- [http://www.dmtf.org/newsroom/pr/view?item\\_key=3b542cbc5e6fc9ede97b9336c29f4c342c02c4e9](http://www.dmtf.org/newsroom/pr/view?item_key=3b542cbc5e6fc9ede97b9336c29f4c342c02c4e9)
- <http://www.forrester.com/Research/Document/Excerpt/0,7211,42968,00.html>

## RELATED ARTICLES:

- A virtual appliance primer: <http://www.linux.com/feature/138166>

**By: Amit Kumar Saha** is passionate about writing and blogging, and works at Sun Microsystems, India. He blogs at <http://blogs.sun.com/amitsaha>, and can be reached at [amitsaha.in@gmail.com](mailto:amitsaha.in@gmail.com)



# Setting Up A Virtual Infrastructure

Set up an integrated network of virtual machines running in QEMU and UML.



**V**irtualisation is quite hot these days. Companies are trying to cut down costs and reduce huge electricity bills by consolidating servers. A major share of the market in this area belongs to VMware products. Recently, Red Hat and Novell have also entered the field with their Xen-based products. With the acquisition of Innotek (the company that makes VirtualBox) by

Sun Microsystems, the marketing war has intensified. Specifically after Sun open-sourced its Solaris operating system, since the company had already exposed the corporate world to Solaris containers (also known as Zones). I have practically set up and watched Zones in action at a few large companies. But in this "War of the Worlds", people often neglect a powerful open source product called QEMU.

QEMU is a virtual machine program. It runs

on top of an operating system known as the host. Within QEMU, another operating system runs which is known as the guest. QEMU virtualises a complete hardware environment, including the CPU (i386, x86\_64, PPC, Sparc, MIPS, ARM platforms); therefore, it can be used as a machine emulator and a virtualiser. As a machine emulator it can be used to run OSs and programs created for one machine on another machine.

In order to do this, QEMU uses dynamic translation. As a virtualiser, the guest code is executed directly on the host CPU. The virtualiser mode requires the host and the guest to use the x86 architecture and this is possible with an accelerator module, using KQEMU.

Although I have not used or seen anybody using QEMU in production, for my personal testing I often resort to a virtual network of virtual machines using QEMU. I have successfully set up an integrated network of virtual machines running in QEMU, VirtualBox and UML. To achieve this I take the help of a third party program called VDE (Virtual Distributed Ethernet).

This series of articles will describe the detailed set up of a virtual infrastructure. I leave it to you to try it out in a production environment.

## The host

The host is the machine where all the VMs will run. Depending on your requirements, you can have multiple hosts interconnected to each other to form a large environment. I have used my laptop to set up this infrastructure. So all the server and client setting up described in this paper runs from my somewhat old laptop—an IBM Thinkpad R50e with Pentium M 1.76 GHz processor, 1GB RAM, 60 GB hard disk, Ethernet and wireless interfaces. The operating system is Ubuntu Hardy. The server services I run on this laptop are DHCP, DNS, NIS and NFS.

Since I am limited by physical resources, the virtual machines that I create will be light weight, but you can go for heavier configurations if you have the available resources. Even these server services don't need to be running on the laptop—you can run them on other physical/virtual machines if you are an experienced administrator.

If you are following this article in order to learn how various services are set up in the real world, I would advise you to start setting up everything on one machine and move services gradually away to different machines as and when your confidence levels increase.

## The guest

The guests are various virtual machines running under QEMU. All other services will run on the guest operating systems. I have used a variety of guest operating systems such as all three BSDs, various Linux distributions and even Microsoft Windows.

## Setting up the host system

We will start by installing QEMU, VDE and DNSMASQ (our DHCP and DNS server), and then gradually progress towards setting up the other server services.

### Installing and setting up QEMU and VDE Switch

Install the basic essential packages as follows:

```
sudo apt-get install qemu kqemu-modules-2.6.24-19-generic \
vde2 uml-utilities bridge-utils
```

Next, create a *udev* rule for the *kqemu* module so that when the *kqemu* module is loaded it has the correct group ownerships and permissions:

```
sudo sh -c `echo KERNEL=="kqemu", NAME="%k", GROUP="kqemu", \
MODE="0660" > /etc/udev/rules.d/60-kqemu.rules`
```

Now create the modules file for *kqemu* and make sure it is loaded at boot time:

```
sudo sh -c `echo options kqemu major=0 > \
/etc/modprobe.d/kqemu` \
sudo sh -c `echo kqemu >> /etc/modules`
```

Load the tun/tap driver and KQEMU module:

```
sudo modprobe tun
sudo modprobe kqemu
```

We need a tap device for use by VDE Switch, which can be created as follows:

```
sudo tunctl -t tap0 -u ajitabhp
```

The *-u* option will create the interface for a particular user. Now, start the VDE Switch daemon and fix permissions on the control file:

```
sudo /usr/bin/vde_switch -tap tap0 -daemon
sudo chmod -R a+rwx /var/run/vde.ctl
```

At this stage we are ready to create and launch the QEMU virtual machines.

I also want to bring all virtual machines from various virtualisation environments under a single network. Using the VDE Switch and a single tap device we can bring all the virtual machines built using QEMU under a common network. Later, you will see that we will also bring the VirtualBox and UML machines under the same network. To achieve this we need to create a bridge interface and add the tap0 interface to it.

From this point onwards, if you decide to go ahead without using the bridge interface and stick to the virtual network of QEMU virtual machines, then you can proceed as follows:

```
sudo ifconfig tap0 10.111.111.254 netmask 255.255.255.0
```

This tap0 interface will act as a default gateway for this subnet. Now let's take on the IP forwarding and masquerading set-up. You will also need to put tap0 in place of br0 in the DNSMASQ configuration.

I am assuming that you have decided to stick with the bridge interface. Go through the following commands to create the bridge interface:

```
sudo brctl addbr br0
sudo ifconfig tap0 0.0.0.0 promisc
sudo brctl addif br0 tap0
sudo ifconfig br0 10.111.111.254 netmask 255.255.255.0 up
```

To allow packets in the virtual machines to reach the outside world, we need to enable forwarding on the host machine:

```
sudo sh -c `echo 1 > /proc/sys/net/ipv4/ip_forward`
sudo iptables -t nat -A POSTROUTING -o eth1 -j MASQUERADE
```

Finally, for better performance of the QEMU:

```
sudo sh -c `echo 1024 > /proc/sys/dev/rtc/max-user-freq`
```

## Installing and setting up DNSMASQ

DNSMASQ is a lightweight DNS, TFTP and DHCP server. It can provide coupled DNS and DHCP services to a LAN. To install it in Ubuntu, use the following command:

```
sudo apt-get install dnsmasq
```

Next, make the back-up of the existing configuration file before creating a new one as follows:

```
sudo mv /etc/dnsmasq.conf /etc/dnsmasq.conf.bak
```

The following is my DNSMASQ configuration file:

```
domain-needed
bogus-priv
interface=br0
no-dhcp-interface=eth0
no-dhcp-interface=eth1
domain=virtual.lan
dhcp-range=10.111.111.110,10.111.111.115,255.255.255.0,10.255
.255.255.12h
dhcp-host=52:54:00:00:EE:01,openbsd41,10.111.111.101,45m
dhcp-host=52:54:00:00:EE:02,freebsd63,10.111.111.102,45m
dhcp-host=52:54:00:00:EE:03,netbsd70,10.111.111.103,45m
dhcp-host=52:54:00:00:EE:04,plan9,10.111.111.104,45m
dhcp-host=52:54:00:00:EE:05,centos5,10.111.111.105,45m
dhcp-host=52:54:00:00:EE:06,opensolaris,10.111.111.106,45m
```

You can change a few options as per your requirement.

Please check the man page to know what each option does. I prefer to keep a static lease for a few of my operating systems, but there is no requirement of this at all. Also, you can change the domain name and DHCP range as per your requirements. I have specified that DNSMASQ should listen on the br0. You can change it to tap0 if you have decided not to continue with the bridge method. DNSMASQ automatically adds the loopback (local) interface to the list of interfaces to use with this option.

Of course, we need to start the DNSMASQ service as well. The procedure is as follows:

```
sudo invoke-rc.d dnsmasq start
```

Now that we have installed DNSMASQ on this machine, let us take advantage of it to provide the DNS cache for faster browsing. Please note that this is not required to achieve the objectives of this article.

To use DNSMASQ as the local DNS cache, uncomment the following line in */etc/dhcp3/dhclient.conf* and add the IP addresses for the OpenDNS servers (I use OpenDNS, you don't have to):

```
# prepend domain-name-servers 127.0.0.1
```

...so that it looks like the following code:

```
prepend domain-name-servers 127.0.0.1,208.67.222.222,208.67.220.220;
```

Now we will install the virtual machines. Create the disks as follows:

```
dd if=/dev/zero of=openbsd41.img bs=1k count=0 seek=1000000
dd if=/dev/zero of=freebsd63.img bs=1k count=0 seek=1000000
dd if=/dev/zero of=centos5.img bs=1k count=0 seek=2000000
```

I am assuming that you have the ISO files for the installation disks. *vdeqemu* (or *vdeq*) is a wrapper program to start the QEMU virtual machine connected to a VDE network. It uses the *qemu -tun-fd* parameter to set up the connection with a *vde\_switch*.

```
vdeqemu -net vde,vlan=0 -net nic,vlan=0,macaddr=52:54:00:00:EE:01 \
-m 128 -localtime -soundhw all -hda openbsd41.img \
-no-kqemu -cdrom ~/ISO/cd41.iso -boot d
```

```
vdeqemu -net vde,vlan=0 -net nic,vlan=0,macaddr=52:54:00:00:EE:02 \
-m 128 -localtime -soundhw all -hda freebsd63.img \
-cdrom ~/ISO/fbsd63.iso -boot d
```

```
vdeqemu -net vde,vlan=0 -net nic,vlan=0,macaddr=52:54:00:00:EE:05 \
-m 256 -localtime -soundhw all -hda centos5.img \
-cdrom ~/ISO/centos5.iso -boot d
```

As you can see, I have used certain MAC addresses for my set-ups. You can change these as per your

requirements. However, if you do change them, make sure that the change reflects in the DNSMASQ configuration file as well.

During the installation, choose the DHCP mode for network configuration. After the installation is over, start the virtual machines as follows:

```
vdeqemu -net vde,vlan=0 -net nic,vlan=0,macaddr=52:54:00:00:EE:01 \
-m 128 -localtime -soundhw all -hda openbsd41.img \
-no-kqemu 1>/tmp/openbsd41.log 2>&1 &

vdeqemu -net vde,vlan=0 -net nic,vlan=0,macaddr=52:54:00:00:EE:02 \
-m 128 -localtime -soundhw all -hda freebsd63.img \
1>/tmp/freebsd63.log 2>&1 &

vdeqemu -net vde,vlan=0 -net nic,vlan=0,macaddr=52:54:00:00:EE:05 \
-m 256 -localtime -soundhw all -hda centos5.img \
1>/tmp/centos5.log 2>&1 &
```

You can first try by pinging the default gateway 10.111.111.254 and then the other virtual machines. Once you are sure that the basic networking is working between the Vms, as well as the VMs and the host, you can proceed with setting up services. The next section deals with how you can set up a centralised login using NIS and NFS shares between the VMs.

## Setting up portmap, NIS and NFS servers

In order to install the portmap, NIS and NFS servers, use the following commands:

```
sudo apt-get install portmap nis nfs-kernel-server
```

We will configure the NIS master server in the host by changing the *NISERVER* line in */etc/default/nis* file as follows:

```
NISERVER=master
```

In order to set the NIS domain name of the server, we can use the *domainname* command as follows which is a part of the *yp-tools* package:

```
sudo domainname virtual.lan
```

The name of the NIS domain resides in the */etc/defaultdomain* file. You can directly edit this file and put the NIS domain name here also.

Edit the */etc/ypserv.securenets* file and add the IP addresses of the NIS client/slave servers to this file, and don't forget to comment out the entry giving access to the world. This entry will look like what's shown below:

```
0.0.0.0          0.0.0.0
```

After these changes, my */etc/ypserv.securenets* file

looks like what follows:

```
# Always allow access for localhost
255.0.0.0      127.0.0.0

255.255.255.0  10.111.111.0
255.255.255.255 192.168.2.178 # This is the IP of my eth1
interface on the host
```

Now we can add the */var/yp/Makefile* if required. I have edited this file and made a few changes. I have changed the minimum UID and minimum GID that will be included in the *passwd* and *group* maps.

The default entry in my Makefile was 1000 for both of them, which were my user ID and default group IDs and I don't want to make myself a part of NIS maps. So I changed these values to 1001. The entries now look like what follows:

```
MINUID=1001
MINGID=1001
```

Note that I've not listed my entire Makefile because it's too large.

Following this I set the option for merging the *passwd* and *shadow* files because I am planning a mix of operating systems, some of which do not support shadow files:

```
MERGE_PASSWD=true
```

Next, comment out the line containing the source file location of *AMD\_HOME*. The reason I am commenting this line is because I will be generating the AMD map on-the-fly from the AUTOFS map definitions.

```
#AMD_HOME      = $(YPSRCDIR)/am-utils/amd.home
```

Now since I only want few maps to be built, following are my changes:

```
ALL = passwd group hosts rpc services netid protocols netgrp
#ALL += publickey mail ethers bootparams printcap
#ALL += auto.local
ALL += amd.home amd.master auto.master auto.home
#ALL += timezone locale networks netmasks
```

As written above, since I want the AMD maps built on the basis of the AUTOFS maps, I have to comment out the section of the Makefile which generates the AMD map as shown below:

```
#amd.home: $(AMD_HOME) $(YPDIR)/Makefile
#      @echo "Updating $@"
#      -@sed -e "s/#.*$/d" -e "/^$/d" $(AMD_HOME) | \
#      $(AWK) '{\`
```

```

#         for (i = 1; i <= NF; i++) \
#             if (i == NF) { \
#                 if (substr($$i, length($$i), 1) == "\") \
#                     printf("%s", substr($$i, 1, length($$i) - 1)); \
#                 else \
#                     printf("%s\n", $$i); \
#             } \
#             else \
#                 printf("%s ", $$i); \
#         }' | $DBLOAD -i $AMD_HOME -o \
$(YPMAPDIR)@0 - $@ \
#         -@$(NOPUSH) || $(YPPUSH) -d $(DOMAIN) $@
```

Finally, add the following lines at the bottom of the Makefile to generate the AMD maps from the respective AUTOFS definitions:

```

amd.home: $(AUTO_HOME)
    -@if [ -f $(AUTO_HOME) ]; then \
    sed -e "/^#/d" -e s/#.*$// $$(AUTO_HOME) \
    | $(AUTO2AMD) | $DBLOAD -i amd.home -o $(YPMAPDIR) \
    amd.home \
        - $(YPDIR)/$(LOCALDOMAIN)/amd.home; \
    echo "updated amd.home"; \
    if [ ! $(NOPUSH) ]; then \
        $(YPPUSH) amd.home; \
        echo "pushed amd.home"; \
    else \
        :; \
    fi \
    else \
        echo "couldn't find $(AUTO_HOME)"; \
    fi

amd.master: $(AUTO_MASTER)
    -@if [ -f $(AUTO_MASTER) ]; then \
    sed -e "/^#/d" -e s/#.*$// -e s/auto/ amd/ $(AUTO_ \
MASTER) \
    | awk '{ printf("%s %s\n", $$1, $$2 ) }' \
    | $DBLOAD -i amd.master -o $(YPMAPDIR)/amd.master \
        - $(YPDIR)/$(LOCALDOMAIN)@0; \
    echo "updated amd.master"; \
    if [ ! $(NOPUSH) ]; then \
        $(YPPUSH) amd.master; \
        echo "pushed amd.master"; \
    else \
        :; \
    fi \
    else \
        echo "couldn't find $(AUTO_MASTER)"; \
    fi
```

The *amd.home* NIS map will be generated from the *auto.home* map, while the *amd.master* map will be generated from the *auto.master* map.



Different flavours of UNIX use different auto-mounting mechanisms to mount the network shares automatically, on demand. Linux uses *autofs* and the BSD uses *AMD* (auto mount daemon). Since the virtual machines are different flavours of BSD and Linux, I have to support both these types of map definitions, and to avoid duplication of definitions I have taken the help of Makefile to generate the AMD maps on the basis of *autofs* maps.

Finally, it's time to build the NIS database for the first time:

```
sudo /usr/lib/yp/ypinit -mp
```

After changing any of the NIS maps files, make sure to execute the following command:

```
sudo make -C /var/yp
```

The NIS servers can be defined in the */etc/yp.conf* file if they cannot be found by the default method of broadcasting on the local network. The format of the */etc/yp.conf* file is as follows:

```
domain <domainname> server <servername>
```

## Lockdown portmap

Add the following line to the */etc/hosts.deny* file and block all hosts from accessing *portmap*:

```
portmap mountd nfsd statd lockd rquotad ypserv ypbind: ALL
```

Now, add the list of hosts allowed to access *portmap* in */etc/hosts.allow*:

```
portmap mountd nfsd statd lockd rquotad ypserv ypbind: \
10.111.111. 192.168.2.178
```

Although *tcp* wrappers allow hostnames, because of a limitation in *portmap*, the list of allowed hosts have to be IP addresses.

## Create NFS shares

Edit the */etc/exports* file and add the following NFS share:

```
/export/src 10.111.111.0/24(rw,subtree_check,no_root_squash)
```

Note that this is just a test share as of now.

Now, export the shares by executing the following command:

```
sudo exportfs -ra
```

This command has to be executed whenever the */etc/exports* file is modified.

## Setting up auto-mounter

Auto-mounter can be used to mount NFS shares, such as home directories of users, as and when they are accessed. In my set-up, *autofs* is not used on the host machine, so I don't need *autofs* on this machine. But if required, it can be installed as follows:

```
sudo apt-get install autofs
```

To set up auto-mount for home directories, create the */etc/auto.master* file and add the following line to it:

```
/home    /etc/auto.home
```

Next, create the */etc/auto.home* file and add the following line to it:

```
*      gateway.virtual.lan:/export/home/&
```

After changing the auto-mounter maps we have to regenerate the NIS database; so, run the following command:

```
sudo make -C /var/yp
```

## Create an NIS user

Let us now create an NIS user and group. The home directory of the user will be mounted automatically on whatever server the user logs into:

```
groupadd -g 1001 nisusers
useradd -u 1001 -g users -m -d /export/home/nisuser01 -s
/bin/sh -c "Test NIS User" nisuser01
usermod -d /home nisuser01
passwd nisuser01
```

The home directories of all users in NIS are located in */export/home*, but whenever the user logs in, it will be auto-mounted via NFS under */home*, as specified in the *autofs.home* and *amd.home* NIS maps. For this reason, after the user is created we will change the location of its home directory to */home* so that the *passwd* file is updated with the correct path.

Creating a user changes the *passwd* NIS map. So, we need to regenerate the NIS maps as follows:

```
sudo make -C /var/yp
```

## All set to go! Start the services

Now that all is set to go, let's start the services:

```
sudo invoke-rc.d portmap start
sudo invoke-rc.d nfs-kernel-server start
sudo invoke-rc.d nis start
```

To test whether NIS is working, we can quickly run a

few commands on the host. To see the name of the NIS server that the host is bound to, we can use the *ypwhich* command. Even if you see something like *localhost.localdomain*, it's fine, as the name is being picked from the */etc/hosts* file. But on the VMs, those commands should show you *gateway.virtual.lan*.

To quickly view the contents of the exported *passwd* map, execute the following:

```
ypcat passwd
```

It should show you the entry for the *nisuser01*, which we created.

Note that the authentication of our host is not through NIS, and I would recommend that you leave it that way.

## Setting up the guest systems

For the guest systems, I used three OSs: CentOS, OpenBSD and FreeBSD. Let's take them one by one.

### CentOS 5

In CentOS (or RHEL and Fedora), to make use of NIS to authenticate users in the system, run the following commands as the root user:

```
domainname virtual.lan
echo "virtual.lan">>/etc/domainname

authconfig --enablenis --nisdomain=virtual.lan \
--nisserver=gateway.virtual.lan --updateall
```

After this, log on to the virtual machine as the root and then test the NIS using the following code:

```
ypwhich
```

```
ypcat passwd
```

If all works as expected, then you can log out and log in as *nisuser01*; you will see that the home directory is automatically mounted for this user.

### OpenBSD

In the OpenBSD guest, make the following changes in */etc/rc.conf.local*

```
# Enable RPC
portmap=YES

# Enable Amd automounter
amd=YES

# Enable lockd for NFS locking
rpc.lockd=YES

# Activate the auto mounter
echo "/home amd.home" >> /etc/ amd/master
```

Now, enable the NIS client as follows:

```
# echo virtual.lan > /etc/defaultdomain
```

Usually, the available NIS server is picked by the NIS client using network broadcast, but somehow when I started the *ypbind* service, the client was not able to bind with the NIS server. So I had to specify the NIS server as follows, executing the commands as root:

```
mkdir /etc/yp
echo "gateway.virtual.lan" > /etc/yp/virtual.lan
```

Next append NIS maps to *passwd* and *group* files:

```
# vipw  # Append the following line
+:+::::::::::

# pwd_mkdb /etc/master.passwd

# vi /etc/group  # Append the following line
+:+:::
```

You may also want to prepend a '+' to the existing group entries so that they are overridden by the NIS map entry. For example:

```
+wheel:+:0:root
```

## FreeBSD

To activate NIS and NFS clients in FreeBSD, make the following additions in the */etc/rc.conf* file:

```
nisdomainname="virtual.lan"
nfs_client_enable="YES"
nis_client_enable="YES"
nis_client_flags="-m -S virtual.lan,gateway.virtual.lan"
rpc_lockd_enable="YES"
rpc_stated_enable="YES"
rpcbind_enable="YES"
amd_enable="YES"
amd_map_program="ypcat -k amd.master"
```

Next append NIS maps to *passwd* and *group* files:

```
# vipw  # Append the following line
+:+::::::::::

# pwd_mkdb /etc/master.passwd

# vi /etc/group  # Append the following line
+:+:::
```

You may also want to prepend a '+' to existing group entries so that they are overridden by the NIS map entry. For example:

```
+wheel:+:0:root
```

## Summing up

By using simple NIS and NFS set-ups we have already laid the foundations for a somewhat complex set-up. Now, you can use your imagination to take this forward. All the best! 

## ACKNOWLEDGEMENTS AND REFERENCES:

To set up this virtual infrastructure, I have referred to lots of useful resources on the Internet, without which it would have been very difficult and time consuming to achieve the result. These are listed below.

### Homepages of the tools used in this set-up:

- DNSMASQ home page: [www.thekelleys.org.uk/dnsmasq/doc.html](http://www.thekelleys.org.uk/dnsmasq/doc.html)
- QEMU and KQEMU home page: [bellard.org/qemu](http://bellard.org/qemu)
- VDE home page: [vde.sourceforge.net](http://vde.sourceforge.net)

### Excellent technical references:

- Virtual Square wiki: [wiki.virtualsquare.org](http://wiki.virtualsquare.org)
- QEMU: [alien.slackbook.org/dokuwiki/doku.php?id=slackware:qemu](http://alien.slackbook.org/dokuwiki/doku.php?id=slackware:qemu)
- VDE: [alien.slackbook.org/dokuwiki/doku.php?id=slackware:vde](http://alien.slackbook.org/dokuwiki/doku.php?id=slackware:vde)
- DNSMASQ: [help.ubuntu.com/community/Dnsmasq](http://help.ubuntu.com/community/Dnsmasq)
- NIS: [help.ubuntu.com/community/SettingUpNISHowTo](http://help.ubuntu.com/community/SettingUpNISHowTo); [www.yolinux.com/HOWTO/NIS-HOWTO.html](http://www.yolinux.com/HOWTO/NIS-HOWTO.html); [www.linux-nis.org/nis-howto/HOWTO](http://www.linux-nis.org/nis-howto/HOWTO); [www.linuxhomenetworking.com/wiki/index.php/Quick\\_HOWTO:\\_Ch30:\\_Configuring\\_NIS](http://www.linuxhomenetworking.com/wiki/index.php/Quick_HOWTO:_Ch30:_Configuring_NIS)
- NFS: [help.ubuntu.com/community/SettingUpNFSHowTo](http://help.ubuntu.com/community/SettingUpNFSHowTo); [www.ubuntugeek.com/nfs-server-and-client-configuration-in-ubuntu.html](http://www.ubuntugeek.com/nfs-server-and-client-configuration-in-ubuntu.html); [www.onlamp.com/pub/a/bsd/2002/02/14/Big\\_Scary\\_Demons.html](http://www.onlamp.com/pub/a/bsd/2002/02/14/Big_Scary_Demons.html); [www.freebsd.org/doc/en/books/handbook/network-nfs.html](http://www.freebsd.org/doc/en/books/handbook/network-nfs.html)

### Troubleshooting:

- NFS: [www.ussg.iu.edu/usail/network/nfs/tips.html](http://www.ussg.iu.edu/usail/network/nfs/tips.html); [stromberg.dnsalias.org/~strombrg/NFS-troubleshooting-2.html](http://stromberg.dnsalias.org/~strombrg/NFS-troubleshooting-2.html); [tldp.org/HOWTO/NFS-HOWTO/troubleshooting.html](http://tldp.org/HOWTO/NFS-HOWTO/troubleshooting.html)
- NIS: [www.linux-nis.org/nis-howto/HOWTO\\_troubleshooting.html](http://www.linux-nis.org/nis-howto/HOWTO_troubleshooting.html); [uw713doc.sco.com/en/NET\\_nis/nisN.troub.html](http://uw713doc.sco.com/en/NET_nis/nisN.troub.html); [www.softpanorama.org/Net\\_Application\\_layer/NIS/nis\\_troubleshooting.shtml](http://www.softpanorama.org/Net_Application_layer/NIS/nis_troubleshooting.shtml)

**By: Ajitabh Pandey.** The author has more than 12 years of diversified IT industry experience in training, support and consulting. His website is at <http://ajitabhpandey.info> and you can reach him at [ajitabhpandey@ajitabhpandey.info](mailto:ajitabhpandey@ajitabhpandey.info)





# Virtualisation in OpenSolaris, Part 1



## Zones and ZFS

Virtualisation today is a major mover and shaker in the computing world and every OS vendor is scrambling to provide virtualisation capabilities in this hot and competitive domain. Sun Microsystems, entering the fray with virtualisation capabilities in OpenSolaris, is no exception. In this article we will look at the virtualisation features in OpenSolaris—Zones and ZFS.

In general, virtualisation applies to techniques of abstracting and sharing physical resources using software or hardware. It is the software techniques that are of primary interest today. Virtualisation in OpenSolaris can have different levels, depending on what resources are being abstracted. Each kind of virtualisation has a cost or overhead associated with it, depending on how much resources are being virtualised. The following are the levels of virtualisation in OpenSolaris in order of increasing overhead:

- **chroot:** This is the most basic form of virtualisation creating isolated filesystem subtrees. This has been present in every UNIX/Linux variant and computer science

textbooks since time immemorial.

- **Zones and Resource Controls:** This is a very lightweight form of OS level virtualisation, where the OS creates an illusion of multiple instances of itself running concurrently and making each appear to be an independent machine. This is based on the same basic concept as FreeBSD Jails. Zones have less than 1 per cent overhead.
- **Storage virtualisation:** The ZFS filesystem in OpenSolaris provides a means of abstracting details of physical disks/arrays by providing a uniform storage pool concept. ZFS integrates the volume manager and filesystem in one logical whole and reduces overhead,

providing simplified management and high performance.

- **Network virtualisation and Resource Control:**

The new Crossbow project that is soon to integrate into OpenSolaris provides advanced network virtualisation capabilities, with virtual NICs (network interface cards), flow-based resource control, virtual routers, firewalls and so on. Using simple concepts, it brings in high-end performance and networking capabilities to the OpenSolaris networking stack.

- **Full desktop virtualisation via VirtualBox:**

VirtualBox is a popular open source project that provides full system virtualisation in an easy-to-use product. It is mainly useful in the desktop and workstation domain providing virtualisation for a variety of developer and end-user needs. VirtualBox is available on a large range of platforms including OpenSolaris.

- **Xen hypervisor:** The Xen hypervisor has been ported to OpenSolaris and provides both full virtualisation in HVM mode, as well as paravirtualisation. Xen is a hypervisor where the core virtualisation software runs in a privileged context outside the OS kernel and provides resource arbitration and a messaging bus. In case of paravirtualisation, the guest OS is aware that it is running inside Xen and uses device drivers optimised for Xen.

- **Hardware partitions:** This is a purely hardware-level virtualisation or resource isolation that allows you to create electrically isolated partitions inside a single machine. An example of this is hardware domains on SUN Fire series of SPARC boxes. This technology allocates dedicated resources to each partition, as opposed to sharing. This aspect, however, is beyond the scope of this article.

Apart from resource management and sharing, virtualisation also improves resource utilisation without compromising on safety and stability by allowing more isolated tasks on the same machine to better exploit the machine's resources without stepping on each other. In this article we will be discussing Zones and ZFS.

## OpenSolaris Zones

At the fundamental level, Zones are lightweight sandboxes within an operating system instance in which one or more applications may be installed and executed without an impact or interaction on the rest of the system. The isolation extends throughout the complete OS namespace, resources, and is also secure. This is similar to the Jails facility in FreeBSD and VServer in Linux. There is only one underlying OS kernel, enhanced to provide increased isolation between groups of processes. The Zones facility, in fact, creates isolated containers that appear to behave as an independent OS instance. The benefit of having a single underlying OS is ease of administration and minimal overhead.

The Zones facility introduces a couple of terminologies. The base OS is referred to as the global zone, and the isolated containers are referred to as non-global zones.

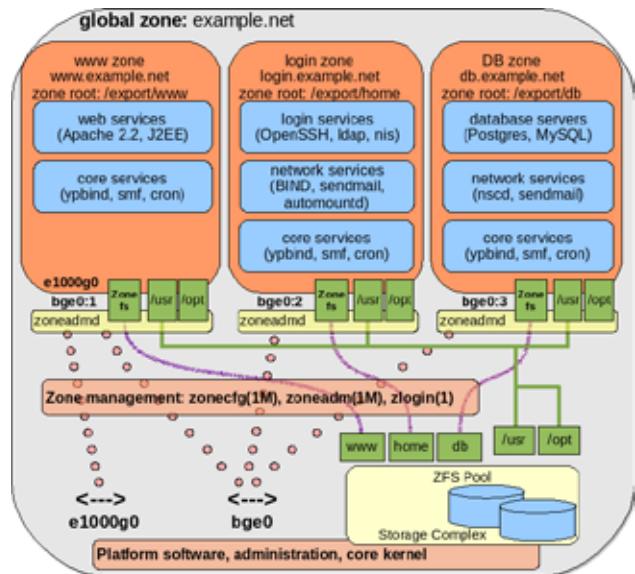


Figure 1: A typical Zones block diagram

A process running in a non-global zone has no access to processes running in the global or other non-global zones. Each non-global zone can be given a logical or dedicated physical network interface, and cannot observe network traffic going to other zones. In addition, each zone is provided a separate root filesystem tree rooted at a subtree of the main filesystem tree akin to *chroot*.

With recent enhancements going on in this space, each one today gets its own TCP/IP stack as opposed to a shared stack. This is known as IP instances. The Crossbow project described later in the article (which will soon integrate into OpenSolaris) introduces virtual NICs and flow control, thereby enabling dedicated end-to-end resource managed network stacks per zone. It is interesting to note that the BeleniX distro website [[www.belenix.org](http://www.belenix.org)] effectively uses Zones and ZFS to create dedicated environments: www Zone for Drupal and Apache2; DB Zone for MySQL; and a *pkg* zone for the package repository and Apache2. There is also an SCM Zone for hosting an internal SVN repository not accessible from outside.

The Zones diagram in Figure 1 tries to comprehensively depict a typical Zone configuration on a server. The three zones provide differentiated services and are sitting on separate ZFS filesystems. The 'e1000g0' NIC lies on the external network and serves the website. The 'bge0' NIC is kept separate from the external network and uses a common means of internal communication between zones as well as for internal login access to the box. The separate login zone is provided for developers to access the system. In addition, the common */usr* and */opt* filesystems are shared from the global zone to reduce disk space requirements.

The Zones framework consists of five principal components:

1. A state model that describes the life cycle of a zone and specifies state transition actions:

Configured <-> Installed <-> Ready  
-> Running -> Shutting Down ->  
Down/Installed.

2. A configuration engine used by administrators to describe and configure a zone, assign resources/resource limits and privileges. The `zonecfg` (1M) command can be used by global Zone administrators to create and configure a zone.
3. Installation support that allows populating the zone's filesystem with the necessary components and also allows patching and upgrading zones.
4. The application environment or 'sandbox' in which processes are run.

In Figure 1, each zone's application environment is depicted by the brown shaded box. This forms the core of Zones implementation. The kernel identifies specific zones via a numeric ID or Zone ID. This ID is reflected in the process structure to cheaply allow segregating Zone-specific processes. Various subsystems like process management, accounting, NFS, IPC, networking, `devfs`, etc, have been given virtualisation capabilities or made Zone-aware. In addition, various commands have also been made Zone-aware. For example, executing `ps -efZ` displays a process's Zone ID. All Zone processes are visible from the global zone. Each zone gets a `zsched` process that is the process 0 (mimic of `sched`) and forms the root of its process tree.

5. The Virtual Platform comprises the set of platform resources dedicated to the Zone. The virtual platform is also responsible for boot, reboot and halt, and is managed by the `zoneadm` daemon. `zoneadm` reads the Zone configuration, creates the central `zone_t` data structure and `zsched` kernel process, sets up virtual network interfaces, and populates devices and the Zone console when starting a zone. It also monitors the zone state as long as the zone is running. The Zone console is created even before the zone is booted so it can mimic a serial console to a physical host.

Each non-global zone, by default, gets less privileges and devices compared to the global zone. Some actions are explicitly disallowed. The only way for zones to communicate with each other is via the networking stack even if they are on the same physical box. In addition, shared filesystems between zones are handled via read-only loopback mounts. To prevent intentional or unintentional cross-zone communication using persistent SYSV IPC objects on read-write shared filesystems, such IPC objects are associated with a zone ID. Thus, such communication is explicitly denied.

In this thread it should be noted that zones can be 'whole root' or 'sparse root'. Whole root zones contain a full copy of the root filesystem hierarchy and are completely

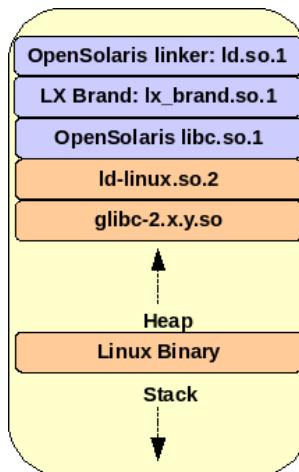


Figure 2: The 'LX' Brand

self-contained. Sparse root zones, as noted in Figure 1, contain some common filesystem sections shared via read-only loopback mounts from the global zone.

## Branded Zones

As a twist to the Zones story, a concept called Branded Zones was introduced. As you can see, Zones allow multiple virtual copies of the 'same' OS environment to be created. Branded Zones extend this by allowing virtual instances of different OS environments to be available. This, however, requires implementation support from the base OS. At present, support is available to create Linux Branded Zones in OpenSolaris. However, it should be noted

that this is just the user environment emulation. There is still a single OpenSolaris kernel running underneath. The Linux kernel does not run.

Branded Zones support adds interposition points in the OpenSolaris kernel: syscall path, process loading, fork, exit, etc. Control is transferred to the brand's kernel support module via these interposition points. This allows a brand to replace or modify default OpenSolaris behaviour and provide alternate behavioural characteristics. The Linux Zone is completely populated with Linux packages and the CentOS distro is officially supported. However, not all Linux kernel functionality is supported like Linux filesystems or device drivers, and not all syscalls are supported. Linux brand support is technically known as the 'LX' brand (Figure 2).

The Linux process loading is handled via the following sequence:

- Kernel jumps into OpenSolaris linker
- OpenSolaris linker loads OpenSolaris `libc` and few other supporting libraries
- The LX brand support library `lx_brand.so.1` is loaded; symbols are resolved by the OpenSolaris linker
- Run `_init()` in `lx_brand.so.1` and pass `lx_handler` address to the kernel
- The handler places extra information needed to exec Linux ELF binaries on the stack in the form of aux vector entries
- Now, jump to Linux linker, which in turn loads `glibc` and others, and resolves symbols in the Linux binaries
- Finally, it jumps into Linux `main()`

Linux system calls are handled via a mechanism called 'trampoline'. The kernel can cheaply identify a so-called branded process since each `proc` structure holds a pointer to the corresponding `brand_ops` vector. OpenSolaris native processes have this field as NULL. So the kernel calls into a brand-specific syscall handler via the `brand_ops` vector. For Linux, this is the `lx_brand` kernel module. This kernel module then transfers control to the userland brand module `lx_brand.so.1` via an upcall.

This is done to keep the implementation simpler and improve stability. The userland brand module performs the necessary emulation and may, in turn, make further native OpenSolaris syscalls. Finally, it returns control directly to the Linux process instead of returning via the same route it was called—trampoline. This eliminates a host of user-kernel transitions and improves performance.

Finally, signals are translated where appropriate, a subset of Linux `/proc` functionality is supported and the necessary devices are provided with the notable exception of a framebuffer device. One advantage of using the Linux brand is that you can use the powerful OpenSolaris debugging features like `mdb` and `Dtrace` directly on Linux binaries.

## Resource Controls and Zones

I have already described Resource Controls/Management on OpenSolaris in detail in the May 2008 issue of LFY. In general, Resource Controls are also a very basic form of virtualisation allowing partitioning and granular allocation of system resources like CPU-binding, thread-limiting, RSS-limiting, CPU-caps, Fair Share Scheduler, etc. All these can be easily applied to Zones, thereby allowing very flexible hierarchical resource management on the system. You can set controlled and granular resource allocation policies and have Zones tuned for differing workloads.

It is also possible for each Zone to have a dedicated, isolated TCP/IP stack with its own routing table, firewall etc. This feature called IP Instances is available in OpenSolaris today. Originally a part of the Crossbow project, this is discussed in the second part of this article on page 51.

## Using Zones

We will now look at a few examples of using the Zones-related commands.

**Creating a zone:** First, let's try some very basic configuration:

```
global# zonecfg -z www
www: No such zone configured
Use 'create' to begin configuring a new zone.
zonecfg:www> create
zonecfg:www> set zonepath=/export/www
zonecfg:www> add net
zonecfg:www:net> set physical=e1000g0
zonecfg:www:net> set address=202.122.64.1
zonecfg:www:net> end
zonecfg:www> exit
global# zoneadm list -vc
```

ID	NAME	STATUS	PATH
BRAND	IP		
0	global	running	/
native	shared		
1	www	configured	/export/www
native	202.122.64.1		

This zone is now in the configured state. Now we can check the zone information. By default, a 'sparse' zone is created. Each zone has a Boolean `autoboot` property that identifies whether that zone should be booted when the system comes up. The SMF service '`svc:/system/zones:default`' is responsible for starting any zone that has '`autoboot=true`'.

```
global# zonecfg -z www info
zonename: www
zonepath: /export/www
autoboot: false.
pool:
limitpriv:
inherit-pkg-dir:
  dir: /lib
inherit-pkg-dir:
  dir: /platform
inherit-pkg-dir:
  dir: /sbin
inherit-pkg-dir:
  dir: /usr
net:
  address: 202.122.64.1
  physical: e1000g0
```

Notice that the zone is a sparse one, since it inherits various top-level filesystem paths from the global zone. A zone can be granted additional privileges identified by the `limitpriv` property. For example, to allow `DTrace` being used inside a local zone, one needs to '`execute set limitpriv=default,dtrace_proc,dtrace_user`' during Zone configuration. Now you can install and boot the zone and log in to its console.

```
global# zoneadm -z www install
A ZFS file system has been created for this zone.
Preparing to install zone <www>.
...
...
global# zoneadm -z www boot
global# zlogin -C www
[Connected to zone 'www' console]
[NOTICE: Zone booting up]
SunOS Release 5.11 ...
...
```

As you can see, `zoneadm` can detect you are using ZFS and auto-creates a ZFS filesystem for the zone. It is easy to create multiple zones of the same configuration by cloning. Since we are using ZFS in this example, cloning a zone is instantaneous via ZFS snapshots.

```
global# zoneadm -z www clone login
WARNING: network address '202.122.61.1' is configured in both
zones.
```

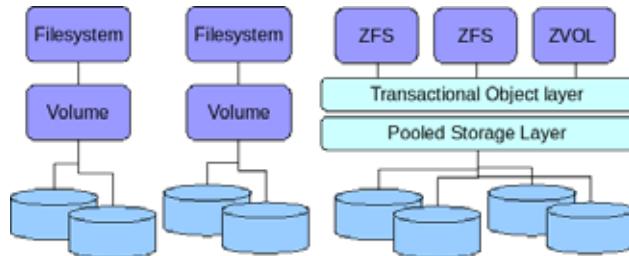


Figure 3: A ZFS pool

```
Cloning snapshot export/www@SUNWzone1
Instead of copying, a ZFS clone has been created for this
zone.
```

You can observe Zone processes from the global zone using `ps -eZ` and `prstat -Z`. The `prstat` command can show per-zone aggregated CPU utilisation. `ifconfig` can show zone information for the network interfaces.

There are many other details that have been obviously glossed over in this article. For further details you can visit the OpenSolaris Zones community portal at [opensolaris.org/os/community/zones](http://opensolaris.org/os/community/zones)

## Storage virtualisation using ZFS

The ZFS filesystem in OpenSolaris is a very innovative ground-up redesign of the traditional UNIX filesystem architecture. ZFS merges ideas from a variety of best practices currently in the market from vendors like NetApp and VERITAS, combined with Sun's in-house R&D and brings to the table a very streamlined and integrated approach to filesystem design. ZFS is a highly scalable 128-bit, object-based, checksummed, transactional, copy-on-write, pooled storage system.

Despite bringing a huge amount of flexibility and features to the table, the focus of ZFS has been on simplicity of configuration and administration. As such, ZFS exports a very straightforward command-line interface that allows the creation of even complex RAID configurations in a few minutes.

It will, of course, take several articles to cover the various ZFS topics and that is not our intention here. Rather, we will take a short look at one critical feature in ZFS that virtualises physical storage without requiring you to go through a complicated circus using Logical Volume Managers.

ZFS introduces the concept of 'pooled storage' that abstracts away the details of the underlying hardware, be it disk arrays or individual disks. The storage hardware is visible as a uniform storage pool. This extremely powerful concept integrated into the filesystem makes storage virtualisation and management a breeze, eliminating the intermediary volume manager altogether. It is easy to grow the storage just by adding more disks to it.

A ZFS pool actually contains one or more ZFS filesystems in it that provide the POSIX filesystem semantics. ZFS filesystem creation is cheap and

instantaneous. In addition, ZFS filesystems can be nested. Each filesystem contains numerous properties that are inherited by the nested ones.

There are multiple ways to configure the ZFS pool. The simplest configuration is a concatenation of the disks providing a sum total storage of all the individual disks. This configuration is not entirely fail-safe since it does not provide redundancy and does not protect against a disk failing. Another configuration is a mirror, where data is replicated identically across all disks in a mirror set. ZFS also provides a variation of RAID-5 called RAIDZ for single-parity RAID and RAIDZ2 for double parity RAID. RAIDZ eliminates the problem of a write hole in RAID5 in which parity can become inconsistent due to a power loss if only a partial stripe has been written. ZFS uses variable stripe widths to eliminate this and can do in software what is not possible with hardware, unless one throws in liberal amounts of battery back-up and expensive NVRAM.

These pool configurations can be combined in interesting ways to give the exact solution one needs. A couple of examples will illustrate how simple it is to create ZFS pools:

1. Create a mirrored ZFS pool with two disks:

```
zpool create foo mirror c0d0p1 c0d0p2
```

2. Attach another device to the mirror. We specify an existing device in the mirror to which another device is to be attached:

```
zpool attach foo c0d0p2 c0d1p1
```

3. Add another mirror group to the existing online pool:

```
zpool add foo mirror c1d0p1 c1d0p2 c1d0p3
```

4. Create a RAIDZ pool with five disks:

```
zpool create bar raidz c1t0d0 c2t0d0 c3t0d0 c4t0d0 c5t0d0
```

In addition, ZFS pools can support separate log devices, hot spares and more recently, supports the use of Solid State Disks (Flash) as an intermediate Level2 cache.

As you can see, ZFS has some amazing properties and allows you ease of use. This short article on ZFS just scratches the surface of how ZFS virtualises storage subsystems -- essentially making expensive hardware RAID redundant. You can refer to the ZFS community portal for more information at [opensolaris.org/os/community/zfs](http://opensolaris.org/os/community/zfs) 

**By: Moinak Ghosh** has over nine years of industry exposure, and was till recently with Sun Microsystems' Solaris Engineering team. He is also the creator of the popular Live CD OpenSolaris distro *BeleniX*.



# Virtualisation in OpenSolaris, Part 2



## Crossbow, VirtualBox and Xen

The previous article on this topic listed a variety of virtualisation techniques in OpenSolaris. We looked at Zones, which is a lightweight OS virtualisation technique, and at storage virtualisation using the ZFS pool concept. Here we will look at network virtualisation and Resource Control using Crossbow, desktop virtualisation using VirtualBox, and the Xen hypervisor.

**V**irtualBox and Xen are fairly established and prevalent multi-platform virtualisation software and have been discussed and covered in various forums and articles. Hence, we will cover the basics on these technologies and how they have been integrated into OpenSolaris. A major portion of this article will be on the exciting new networking infrastructure called Project Crossbow that is currently in beta stage and

will very soon make its debut in the official OpenSolaris release. Since it is an open source project, anyone is welcome to participate. You can check out the project portal at [www.opensolaris.org/os/project/crossbow](http://www.opensolaris.org/os/project/crossbow)

### Crossbow: Network virtualisation and Resource Control

Think about how cool it will be if your OS provided the capability to slice your physical network interface into multiple virtual NICs

with full resource control, network traffic prioritisation and the ability to exploit the hardware flow classification that's now available in commodity network adapters.

Project Crossbow in OpenSolaris makes all these possible by virtualising the network stack and NIC around any service protocol (HTTP, FTP, NFS, etc) or virtual machine. In addition, it introduces a small twist to traditional network packet processing in UNIX that greatly improves performance and scalability.

Each virtual network stack can be assigned its own priority and bandwidth without performance degradation even on a shared NIC. The hardware flow classification engines now being available on commodity network interface cards are exploited to provide high-performance traffic isolation between the virtual network stacks. One point to note here is that project Crossbow is an architectural change in the OpenSolaris networking stack rather than being an add-on layer.

The Crossbow architecture consists of the following functional components:

- **Virtual NICs:** A single physical NIC can be sliced into multiple VNICs and assigned to Zones, Xen or VirtualBox instances. A Virtual NIC is an abstraction that allows identifying and controlling dedicated DMA buffers and Rx/Tx H/W rings. With support from the physical NIC's hardware classifier mentioned below, packets are delivered directly to the IP layer bypassing the data-link layer.
- **Flow management:** A flow is characterised by a class of traffic (protocol, service or virtual machine, etc) and handling policies like bandwidth limit, priority, etc. Flow management and resource control is dependent on a construct called the Serialisation Queue that can dynamically switch the interface from interrupt to polling mode and control the rate of packet arrival from the VNIC. This is the 'twist' alluded to earlier that enables high performance and scalability.
- **Hardware support for flow processing:** Today's modern NICs provide content-based classification and segregation of traffic in hardware, allowing multiple receive and transmit ring buffers per class. These Rx/Tx rings are then associated with flows managed by a Serialisation Queue. This combination of a hardware packet classifier and flow management provides practical network Quality of Service (QoS) in a general-purpose operating system.
- **Administrative interface:** OpenSolaris provides the *dladm* (1M) command to manage link resources. This command is extended to create and manage VNICs. Another command, *flowadm* (1M), can be used to configure flows. In addition, the OpenSolaris Zones administration interfaces are being enriched to seamlessly integrate VNIC and flow management since Crossbow features are a logical adjunct to the Zones facility.

As mentioned in Part 1, the Zones facility uses the

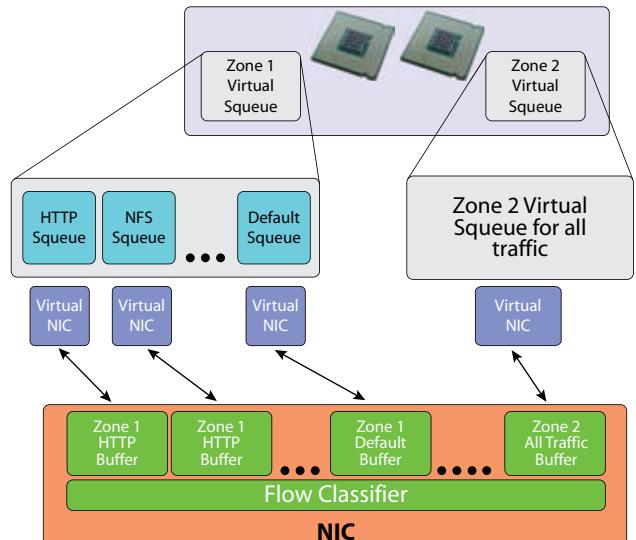


Figure 1: Crossbow functional components

ability to create a dedicated, isolated TCP/IP stack per zone, which is called an IP instance. This feature was implemented as part of the Crossbow project and is already available in OpenSolaris today. The combination of VNICs, flow control and the virtualised TCP/IP stack, makes for a compelling set of possibilities. You can create virtual routers and virtual firewall appliances. You can even create a complete virtual network topology inside a single machine without having to use specialised network simulator software.

The Serialisation Queue with its ability to dynamically switch between interrupt and polling mode is central to the idea of high-performance QOS/Resource Control. Traditional purely interrupt-based packet delivery in general purpose operating systems and host-based QOS has several drawbacks:

1. With layered QOS mechanisms, packets have to be delivered into the system copied into buffers and then post-processed into various queues via the QOS layer. This introduces latency.
2. Since the kernel networking stack is unaware of any bandwidth or other policies, every packet delivered is processed. By the time one decides that the packet needs to be discarded, 80 per cent of the work (interrupt processing, packet copying) is already done. Contrast this to a typical Cisco Router's Fast Switching implementation where Cisco's IOS goes to great lengths to avoid copying packets into the main memory and makes policy decisions as early as possible.
3. Every packet in the host flows through common queues and is processed via common threads making differentiated services difficult. How many times have you been affected by slow interactive terminal response because of someone downloading a large file from the same server?

Moving policy decisions into the networking stack and using Squeues allows overcoming these limitations.

In fact, classification decisions are pushed down as far as the physical network card. Of course, you might ask that since not all network cards are capable of doing hardware classification, what then? Crossbow also includes a software classifier sitting just above the NIC to handle this scenario. Obviously, this degrades performance.

The Squeue controls the rate of packet arrival into the system based on configured bandwidth. Normally, when there are no packets, the NICs Rx ring buffer is in interrupt mode. When the first packet arrives, an interrupt is raised and the packet delivered. At this point the IP layer's Squeue kicks in and switches the Rx ring to poll mode. It then starts pulling in entire packet chains from the NIC as fast as possible. At this point the Rx ring is delivering packets directly into the IP layer, bypassing the data-link layer. Once no more packets are coming in, the Rx ring is switched back to interrupt mode. This might sound a little counter-intuitive to some, but this approach has far less overheads than with the purely interrupt mode, leading to excellent scalability in packet processing. The data-link bypass is not used if the VNIC is placed in promiscuous mode.

The Squeue can take a decision on whether to pull in packets from the NIC or not, thereby effectively limiting bandwidth usage and eliminating wasteful processing. The entire layered architecture uses function pointers abstracted as 'upcall\_func' and 'downcall\_func' to call into each other. In addition, data structures passed between layers are reference counted.

This kind of resource control via Squeues also allows host-based defence against DDOS (distributed denial of service) attacks. Normally, DDOS can cripple all kinds of network services on a machine. With Crossbow, only impacted services segregated by flows, take the hit. Under attack, the priority of the impacted flows can be reduced to allow other network services to continue to function.

Another advantage of this framework is that fine-grained accounting comes for free. It is possible to do per-Squeue accounting to track usage by Virtual Machine, service or protocol. These kernel stats can be queried at intervals from userland to do accounting/billing. Running virtual stacks without any bandwidth or other resource limits can help in capacity planning.

One salient point to note here is that VNICs and flows can be independently configured—that is, it is not necessary to configure a VNIC in order to create a flow.

## Crossbow: Administration interfaces

The Crossbow project introduces two commands, *dladm* and *flowadm*, to manage VNICs and flows. These commands are relatively simple to use. The following examples demonstrate a sample usage:

```
>>> Create a simple VNIC attached to the bge0 interface. The
vnic will be named vnic1.
# dladm create-vnic -d bge0 1
```

```
# ifconfig -a
bge0: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu
1500 index 3
      inet 192.168.1.1 netmask ffff0000 broadcast 192.168.1.255
        ether 0:10:18:a:29:44
vnic1: flags=1000843<UP,BROADCAST,RUNNING,MULTICAST,IPv4> mtu
1500 index 4
      inet 192.168.1.2 netmask ffffff00 broadcast 192.168.1.255
        ether 0:10:18:c:77:55

>>> Now create a VNIC with a guaranteed B/W of 600Mbps &
priority Hi
# dladm create-vnic -d bge0 -m factory -b 600m -G -p high 2

>>> Create a flow for inbound HTTPS traffic so that it gets
dedicated TCP/IP resources
# flowadm add-flow -l bge0 -a transport=TCP local_port=443 https-
1

>>> Modify the flow to add a bandwidth cap, high priority and
the flow processing
>>> bound to CPU# 9
# flowadm set-flowprop -p maxbw=500Mbps,priority=high,cpus=9
https-1
```

Historical flow usage statistics can be recorded by enabling the extended accounting feature of OpenSolaris:

```
>>> Enable extended accounting for network
# acctadm -e extended -f /var/log/net.log net

>>> Later summary information can be displayed
# flowadm show-usage -f /var/log/net.log
Bytes  Packets  Errors  Duration  Bandwidth  Link/Flow
-----+-----+-----+-----+-----+-----+-----+
546908    1031      0     100    43.75 Kbps  flowtcp
0        0        0        0      0.00 bps   flowudp
```

In addition to displaying summary information, detailed date/time-wise information can also be queried. In order to experiment with Crossbow you will need to download a custom OpenSolaris build ISO containing the Crossbow bits from [www.opensolaris.org/os/project/crossbow/snapshots](http://www.opensolaris.org/os/project/crossbow/snapshots).

The *dladm* and *flowadm* man pages are available at [dlc.sun.com/osol/netvirt/downloads/current](http://dlc.sun.com/osol/netvirt/downloads/current)

## Crossbow: Real world use cases

Network virtualisation and flow control is useful in a variety of scenarios:

1. The flow prioritisation and accounting features can be very useful at high-volume brokerage and trading sites serving a large number of customers.
2. Large ISPs want to deploy virtual systems on the same physical machine. This scenario is very common today. ISPs sell each virtual system with differing prices and

QOS guarantees. The ISP can either use a per-VM dedicated physical NIC, which is a more expensive solution, or use Crossbow to more effectively slice and dice shared NICs. This also improves utilisation.

3. Large enterprises commonly use workgroup servers for mundane administrative as well as critical network traffic. A technical support starting a netbackup can impact users doing critical work on the same server. Of course, enterprises have elaborate means to pre-announce and schedule such work mostly on weekends. With Crossbow, this overhead can be reduced by controlling priority and bandwidth usage of the netbackup such that it can be started at any time without impacting critical usage.

In addition to these, Crossbow also brings in possibilities of creating network appliances. One such project is the Virtual Network Machines project on the OpenSolaris website: [www.opensolaris.org/os/project/vnm](http://www.opensolaris.org/os/project/vnm)

At present there is a usable beta Virtual Network Router available for download. Of course, there are other possibilities like firewalls, load balancers, QOS aware NAS and so on. Visit the above project portal for an interesting perspective.

## Desktop virtualisation using VirtualBox

VirtualBox is a high performance Type-2 hypervisor that uses a combination of virtualisation techniques to run many different unmodified operating systems in x86 virtual machines. It is a Type-2 hypervisor since it is hosted inside an operating system. It is highly portable across different host operating systems and supports a wide range of guest operating systems.

VirtualBox uses the client-server architecture with a back-end VBoxSVC process keeping track of VMs and their state. The clients provide the front-end user-interface, whether CLI or GUI, and interact with the back-end via an XPCOM API interface. Presently, there is a Qt-based GUI and an SDL-based GUI. There is a CLI called VBoxManage and an RDP server as well.

VirtualBox tries to optimise and run as much of the guest code as possible, natively, on the CPU avoiding emulation overhead. There are certain contexts (like Real-Mode emulation) where slower emulation may be necessary and it uses an x86 emulator based on QEMU. Traditionally, on x86 the operating system kernel runs at Privilege Level ring-0 while the user context runs at Privilege Level ring-3. When running the guest OS kernel, VirtualBox cannot run it at ring-0 since it has to be in control of the guest. So VirtualBox has a host kernel component that runs at ring-0 and the guest OS kernel is run at ring-1 Privilege Level. VirtualBox sets up CPU and memory contexts exactly as the host expects it and fools the guest OS kernel into thinking that it is running at ring-0.

VirtualBox can also do very advanced code scanning to eliminate problems caused by running ring-0 kernel code in ring-1. Most privileged ring-0 instructions in kernel

generate faults in ring-1 and some behave differently! So VirtualBox has an advanced re-compiler that dynamically scans for such instructions at runtime and patches the machine code to replace such instructions with calls into the hypervisor that can safely emulate those. All these techniques improve performance greatly. VirtualBox also has the ability to use the x86 virtualisation extensions AMD-V and Intel VT-x.

Originally developed by Innotek, it was later acquired by Sun. Support for OpenSolaris as a host OS has been evolving from version 1.3 onwards, and today most of the functionality is supported out of the box. In addition, there is work going on to make VirtualBox on the OpenSolaris host work nicely with VNIC support for Crossbow. In fact, the VirtualBox front-ends should be able to provide seamless creation of VNICs per host OS network device configured.

Using VirtualBox on an OpenSolaris host is really easy and works extremely well. Do try it out.

## Xen on OpenSolaris

The Xen project is an open source Type-1 hypervisor developed at the University of Cambridge computer lab. Xen is a Type-1 hypervisor since it runs on bare metal at the highest privilege level. Though it does not run on a host OS, it still depends on a running OS instance called the control domain or dom0.

Xen supports multiple operating system instances running simultaneously on the same machine and arbitrates access to the underlying hardware. Each running instance is called a domain. The actual hypervisor provides the low-level infrastructure to virtualise the hardware for tasks like handling privileged instruction traps providing a communications bus, etc. However, Xen depends heavily on dom0 (control domain) for most of the other tasks like creating additional domains, allocating memory and other resources, device access, etc, since Xen does not include any device drivers.

As most of you may know, Xen supports both paravirtualisation and full virtualisation -- or HVM for short. In paravirtualisation, the guest OS is aware that it is running inside Xen and makes calls into the hypervisor for privileged operations like sending interrupts across CPUs, accessing devices, etc. These are called hypercalls. Paravirtualisation does not require any special hardware support, but requires that the guest OS be modified to support Xen. In full-virtualisation mode, the guest OS is unaware of Xen, so Xen depends on hardware extensions—AMD-V, Intel VT-x, etc—to intercept privileged operations like physical device access. In general, paravirtualisation has less overheads than the full-virtualisation mode.

Xen has been ported onto OpenSolaris and can use OpenSolaris both as dom0 and as a paravirtual guest OS. Note that since dom0 must communicate with the underlying hypervisor, it is always running in paravirtual mode. Xen on OpenSolaris is referred to as xVM.

Xen on OpenSolaris virtualises the CPU, memory and devices. One or more virtual CPUs (VCpus) are assigned to each domain. Each VCPU preserves the physical CPU state (registers, flags, etc). A VCPU is a schedulable entity like threads in OpenSolaris.

OpenSolaris manages memory in pages like so many other operating systems. So there are the traditional virtual-page-to-physical-page translations done in the kernel. The physical page numbers are obtained from the BIOS or firmware when executing the kernel on bare metal. However, in the Xen context, Xen provides a virtualised view of the physical pages. There are no direct relations between virtualised physical page numbers to the real physical pages of the underlying hardware—this is for supporting live migration of domains. Thus Xen maps the guest OS's physical page numbers to the so-called machine frame numbers of the hardware.

Device support is provided with the help from the control domain or dom0. All the 'real' device drivers reside in dom0 and they directly talk to the hardware. However, since dom0 itself is also running on Xen, it is also subject to physical-page-to-machine-page translations. Thus, all the drivers must use the standard OpenSolaris device driver interface (DDI). As a side note, OpenSolaris offers this standard DDI to provide a stable backward-compatible interface to device driver developers. When a fully-virtualised domain makes an I/O access, it is trapped by the hypervisor, which then forwards the request to the appropriate driver in dom0. In a paravirtualized domain, special PV drivers use a 'front-end' and 'back-end' approach. The PV drivers in the domain (or domU) make hypercalls to communicate over the Xen communication bus to the back-end driver in dom0. The back-end drivers then call into the real device drivers in dom0 that talk to hardware. Since paravirtual mode avoids expensive privileged call trapping by the hypervisor, it is much more efficient.

OpenSolaris provides a paravirtual block device and a network device. In addition, a virtual graphics console is also provided in full-virtualisation mode. The paravirtual mode does not yet have a virtual framebuffer device and uses an integrated VNC server instead for graphics output.

The hypervisor code requires quite a bit of modifications to make it work on OpenSolaris. Most of these patches have been pushed into the upstream Xen project. However, there are some patches that are not accepted and need to be maintained separately. So you cannot simply download the Xen source and expect it to build on OpenSolaris.

Patched Xen code is available from the OpenSolaris website corresponding to each biweekly OpenSolaris build that is released. Check here: [dlc.sun.com/osol/on/downloads](http://dlc.sun.com/osol/on/downloads)

Since OpenSolaris supports kernel crashdumps, the presence of Xen provides an interesting situation. If a user domain running OpenSolaris as the guest or the control domain (dom0) crashes, then the kernel crashdump is stored as usual in the virtual disk. That crashdump does not contain any hypervisor information

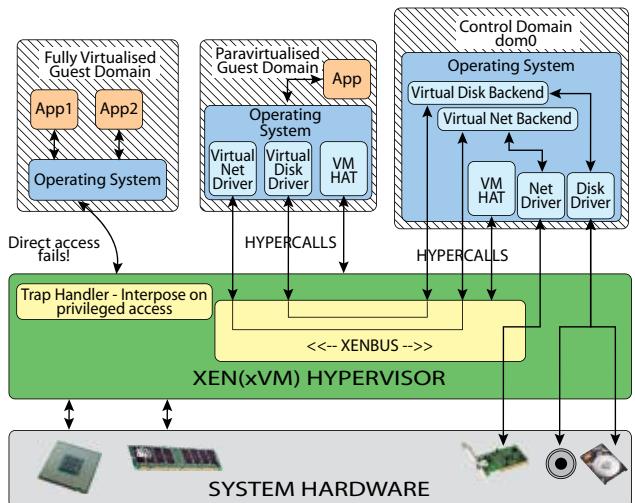


Figure 2: Xen(xVM) Framework

since the hypervisor is off-limits to all domains except via hypercalls. However, if the hypervisor itself crashes, then it will produce a crashdump in dom0 that contains both hypervisor state and dom0 kernel state with symbolic information. In such a kernel core file, the hypervisor appears as a simple kernel module called xpv.

Finally, since the hypervisor is completely isolated even from the dom0, there is no way to trace directly into it. However, one can, of course, trace the hypercalls from dom0. Executing full hypercall tracing using `dtrace -n 'xpv::: /'` while doing Xen operations provides a detailed view of the steps taken for actions like creating a domain, migrating domains, booting them, etc.

You can visit the Xen community portal at [www.opensolaris.org/os/community/xen](http://www.opensolaris.org/os/community/xen) if you are interested to know more and want to play with Xen on OpenSolaris.

## Bootnote

Virtualisation is a vast topic and these two articles give a high-level introduction to virtualisation technologies on OpenSolaris. It is not possible in the limited space of a couple of articles to do justice to these technologies. In addition, there is yet another virtualisation technology on OpenSolaris that was not covered, called Logical Domains (LDOMS), which is available on SPARC hardware. LDOMS is conceptually somewhat similar to Xen, though it depends on support from system firmware. More information on LDOMS can be found at [www.opensolaris.org/os/community/ldoms](http://www.opensolaris.org/os/community/ldoms)

Finally, there are some good presentations on OpenSolaris virtualisation and OpenSolaris, in general, at [opensolaris.org/os/community/advocacy/os-presentations](http://opensolaris.org/os/community/advocacy/os-presentations) 

**By: Moinak Ghosh** has over nine years of industry exposure, and was till recently with Sun Microsystems' Solaris Engineering team. He is also the creator of the popular Live CD OpenSolaris distro *BeleniX*.

# CodeSport

Welcome to another instalment of CodeSport. This week, we discuss a couple of well-known graph algorithmic problems.

Thanks to all the readers who sent in their solutions and comments to the problems we discussed in the previous column. Last month's takeaway problem was on graph matching. You were to imagine yourself as the lecturer of a class and you wanted to pair up students to work on your assignments. You also know that certain students did not work well with certain others. Hence, you wanted to avoid such pairings. Given that you knew which students do not work well together, you were asked to pair up all your students such that both members of the pair work well with each other? If this were not possible, the challenge was to find the maximal number of student pairs you could form.

Though it seems abstract, this is a well-known graph theory problem called finding the maximal matching in a given graph. For our problem, let us model the students as vertices in a graph. A pair of students who can work well together have an edge connecting them. We want to pick up pairings of students such that:

- a) All students are picked up,
- b) We do not pair a student with someone he or she does not work well with.

Such a listing of student pairs is known as a 'matching' in graph theory terms. A matching is a collection of independent edges with disjoint end points, i.e., no two edges in a matching share any end points. Our first requirement of pairing up all students of the class is known as perfect matching in a graph, where all the vertices of the graph are covered by the matching we determine. A perfect matching requires that we cannot leave any student out. Is this always possible? What would happen if your class had an odd number of students? There is no way you can pair up all students since in the end you will be left with a lone student for whom you can not find a mate. Is an odd numbered class the only case where a perfect matching is not possible?

Consider the case of four students, where one student works well with three other students, but none of the other three students can work well with each other. In this case, there is no way to get a perfect matching, because once you

select one edge (one student pair) for your matching, none of the remaining edges are independent of the chosen edge, and hence cannot contribute any more edges to the matching. So perfect matching is not always possible even in the case of graphs that have even a number of vertices. In cases where a perfect matching is not possible, how do we go about finding a maximal matching?

We need some graph theory terminology here. A matched edge is an edge that is included in a matching and an unmatched edge is one that is not included in the matching. Matched vertices are those that are end points of a matched edge and unmatched vertices are those that are end points of an unmatched edge. An alternating path in the graph is a path consisting of alternate matched and unmatched edges. An augmenting path is an alternating path that starts at the unmatched edge and ends at an unmatched edge. Armed with this terminology, let us consider a contrived example, where we have four students—A, B, C and D. A can work well with B and C. D can work well only with C. So we have three edges (A, B), (A, C) and (D, C). Assume that we have chosen a matching consisting of the edge (A, C). Now we cannot extend this matching any more because none of the other edges have disjointed end points with this edge. So we are stuck with a matching of size 1. Is this the maximum-sized matching possible for this graph?

By trial and error, we can find that the maximal matching is by choosing the edges AB and DC and discarding the edge AC. Now we get a matching of size 2. What is the procedure we can follow to find this maximal matching? Assume that we have an existing mapping consisting of the edge AC in the original graph. We find that there is an augmenting path starting from B, traversing the edges BC, AC and CD. Now, if we take a symmetric path difference of the augmenting path and the existing machine, we get a bigger matching. Note that a symmetric path difference of two paths, P1 and P2, consists of the union of those edges in P1 and not in P2, and those edges in P2 and not in P1. The bigger matching consists of the edges AB and DC. The presence of an augmenting path allows us to extend an existing matching

by one more edge, since we now cover the same vertices as before, as well as covering two additional vertices that were unmatched in the earlier matching. Hence, the problem of finding the maximal matching reduces to that of finding augmenting paths.

As long as we can find an augmenting path, we can extend the existing matching by one more edge, taking the symmetric path difference between the newly discovered augmenting path and the existing matching. When we cannot find any more augmenting paths in the graph, we have no way of extending the matching. So we have arrived at the maximal matching possible. Of course, now the question is reduced to that of finding an augmenting path in a graph. For now, I leave it to the reader to come up with a solution. We will discuss the problem of finding augmenting paths in next month's column when we discuss network flows on graphs.

In this month's column, let us look at another well-known graph problem, that of finding the shortest paths from a specified source vertex to all other vertices in a weighted graph. There are well-known algorithms such as Djikstra's Single Source Shortest Path Algorithms (SSSP), for finding the shortest path from a given source to all vertices in a general weighted graph. Djikstra's SSSP algorithm is a greedy one that works by the principle of relaxation, wherein we start with an initial shortest path estimate for a vertex from the source, equal to infinity, and continually relax it until it becomes equal to the actual shortest path distance from the source to that vertex. Each time, we check whether a new path we look at can give a value shorter than the one we already have. If so, we update the shortest path estimate. We maintain the shortest path estimates in an array  $d[i]$  where  $1 \leq i \leq n$ , and 'n' is the number of vertices in the graph. The relaxation operation can be defined on an edge  $(u, v)$  as follows:

```
Relax(u, v, w)
{
    if (d[v] > d[u] + w[u, v])
        d[v] = d[u] + w[u, v];
}
```

All the shortest path algorithms use relaxation as their main operation. But the order in which they relax the graph edges as well as the number of times they relax the edges are different for each algorithm. Each time Djikstra's algorithm picks up the edge with the shortest path estimate to relax first, it relaxes all the edges of the graph  $G$  only once. We have a graph  $G$  with vertices 'n', their edge weights given by the  $w[i, j]$  where  $i$  and  $j$  are the end points of each edge;  $d[]$  is the shortest path estimate array and the source vertex is 's'. Here is the pseudo code for Djikstra's SSSP algorithm.

```
Djikstra(G, W, s)
{
    For (vertex v = 1 to n)
    {
```

```
        d[v] = infinity;
    }
    d[s] = 0;
    Initialise a queue Q to contain all vertices of
    the graph.

    While (Q not empty) //takes O(|V|) time where
    |V| is the number of vertices in 'G'.
    {
        Extract the vertex 'u' from the queue
        which has the minimum d[u] value;
        // we can see that all edges of 'G' are
        examined once, So takes O(|E|) time
        For (each vertex 'v' adjacent to u)
        {
            Relax(u, v, w(u, v));
        }
    }
}
```

What is the running time of Djikstra's SSSP algorithm? We execute the "while Q not empty"  $|V|$  times, where  $|V|$  is the number of vertices in the graph. Hence, we perform the *extract minimum operation*  $|V|$  times. If we use an array to maintain  $d[v]$ , each *extract\_min* operation will take the time of the order of  $|V|$  since we need to search the entire array. We relax each edge only once, so the total time taken for relaxation is of the order of  $|E|$ , where  $|E|$  is the number of edges in 'G'. So the total time for Djikstra's SSSP is  $O(V^2 + E)$ .

How can we improve the running time of Djikstra's algorithm? Can we improve the time taken for *extract\_min* operation? If we use a binary heap instead of an array for maintaining the distance estimates, we can perform *extract\_min* in  $O(\log V)$  instead of  $O(V)$ . Hence, the total time comes down to  $O(V \log V + E \log V)$ . So I leave readers with this question: what data structure can we use to improve this even further?

For this week's takeaway problem, consider a variant of Djikstra's single source shortest path algorithm -- that of finding the shortest paths in directed acyclic graphs (DAG) with no negative weighted edges. Can you come up with an algorithm that can solve the SSSP problem for DAGs in  $O(V+E)$  time complexity?

If you have any favourite programming puzzles that you would like to discuss on this forum, please send them to me. Feel free to send your solutions and feedback to *sandyasm\_AT\_yahoo\_DOT\_com*. Till we meet again next month, happy programming! 

**Sandy Mannarwamy** is a specialist in compiler optimisation and works at Hewlett-Packard India. She has a number of publications and patents to her credit, and her areas of interest include virtualisation technologies and software development tools.



# Spying Eyes Around? Time to Lock Your Data

Come on in, and try encrypting your files. There's nothing to worry about as long as you don't forget your passwords, and anyway, you can sure take regular back-ups, right?

**I**n the early days, when there were no networks to speak of, the field of information security did not exist. Maybe the huge import duties had something to do with it, but the actual computers were considered much more valuable than the (little) information stored on them. Fortunately, sizes were measured in metres and tonnes, so they were quite safe from theft.

Now computers are a commodity, and the information on them is the prize. Not only is there a lot more of it than there used to be, it's a lot more useful. What could anyone do with a file full of bank account details 20 years ago? Nothing much. Now, even the most insignificant personal details can become avenues of attack. Whether it's a home machine or a work machine, there are plenty of things on them that should not fall into the wrong hands. Worse, it is easy

to swipe a USB stick, copy the contents, and put it back without the owner even knowing that it happened.

Clearly, the best solution is to encrypt your files so that they cannot be accessed without a password, and anyone stealing the disk sees only garbage. Unfortunately, there's a very thin line between, "My data is safe from theft," and "Ooops, I can't access my data any more!" And this scares a lot of people away. I'm going to try and make things a little easier, and perhaps give you the confidence to at least try it out on some of your less critical data first. As usual, I'll focus more on the concepts and ideas than on the actual commands.

This article introduces a lot of terminology, so here's a quick overview. 'dm-crypt' is the kernel module that provides transparent encryption of block devices. *dm-crypt* depends upon the 'device mapper', which is a generic framework to map one block device onto another. The userspace tool that makes *dm-crypt* usable is *cryptsetup*. LUKS is an extension to *cryptsetup* that adds even more usability, interoperability, and enhanced security to *dm-crypt*. (LUKS is what this article is mainly about, but it's a complex beast so we sneak up on it slowly!)

Finally, there is EncFS, which is much more lightweight, is based on a completely different mechanism (called FUSE), and can even be used in an ad-hoc manner for those times when you suddenly realise you need encryption for some directory.

## Decisions, decisions, decisions...

There're a lot of choices in the technology available, and it's important to think about what it is you're trying to protect and from whom you're doing so. At one extreme, if you're, say, a RAW agent or a top cop, you might want to make sure that everything is encrypted and nothing is ever kept open at any time.

If you're the James Bond type, on the other hand, you always have a cover story; so you want the system to look normal, with a lot of innocuous files visible, while the real stuff is hidden away so well no one even knows it exists. This is called plausible deniability; we will not be going into this topic either.

At the other end of the scale, most normal people might be satisfied with protecting the actual files on the hard disk or the USB stick. We don't anticipate that our stuff is interesting enough to warrant more than a modest amount of effort. For instance, I do not care if people know that I have about 45 mail folders ranging in size from 1 to 220 MB in size, totalling about 1.2 GB; it is sufficient for me that no one can actually read the e-mails themselves. If I caught someone opening up my PC and freezing the RAM chips, my sense of self-worth would certainly go up dramatically; I might even ask my boss for a big raise!

A lot of current distributions allow you to encrypt the entire disk or just some partitions. It's actually become

## Extreme protection

Encrypting /home is not enough: temporary files are often created in /tmp or /var/tmp, the swap partition often has interesting stuff sitting around, or someone with a Live CD could even tamper with your boot sequence itself to capture your passwords next time you boot! A recent discovery was that, contrary to what we've always believed, the contents of RAM do not get wiped out instantly when power is switched off, but can be retrieved for up to 10 minutes afterwards (more if the chip can be frozen), and this will almost certainly include keys to your encrypted files! To protect against all this, you'd have to encrypt the entire hard disk, including swap, avoid using the 'suspend/resume' feature, and wait for at least 10 minutes after powering it off before letting the machine out of your sight. We will not be covering the technology at this level.

quite painless, at least in the most common cases, like when you're doing a fresh install. I will not be covering this either, since you can easily find instructions for your favourite distribution, including screenshots, on the Web.

A personal note: I often use the 'suspend to disk', a.k.a. 'hibernate', even on my home desktop computer because it can save a lot of power under the right circumstances. It's possible to get this working while using encrypted swap, but it's a little complicated, and I haven't felt the need to go this far yet.

## A bit of history...

The first encryption scheme to be widely used was called *cryptoloop*. Some of you may know the *losetup* command, which allows you to create a block device out of a regular file (for instance, this is how you mount an ISO image on your hard disk if you want to see what's inside). *cryptoloop* simply added encryption capabilities to *losetup*, by adding an option to specify the algorithm to be used. This system had several weaknesses (briefly: known plain text and watermark attacks), and is now deprecated, so we will not describe it further.

The second method that came into existence was called *loop-aes*. This was technically much better than *cryptoloop*. It addressed the security weaknesses of *cryptoloop*, and added a very nice feature: you could have multiple passwords for the same data, and you could change passwords easily without having to copy all the data.

*loop-aes* never made it into the official kernel, and many distributions did not support or include it. You had to download the software and compile it yourself if you wanted to use it. The sad part is that this appeared to be largely due to interpersonal clashes between the kernel maintainers and the folks behind *loop-aes*. If not for those issues this may well have become the default after some time.

## The device mapper

The device mapper is a Linux framework that allows you to create virtual layers of block devices on top of real ones. It was initially created to address various concerns with LVM (Logical Volume Manager) and its interface, but as part of that rethinking, we got a clean, robust, design that separates what the kernel should be doing from what modules should be doing. Device mapper is now used not just for LVM, but for things like RAID and crypto as well, and can be used by any other application that needs to make block data transformations. And of course, these modules (for some reason they're called 'targets') can be stacked on top of each other too, if you wish.

Specifying a device mapper target consists of creating a 'table' and telling the device mapper to use it to create a mapping. This table format isn't very user-friendly, although that shouldn't worry us, as you'll see later. For our purposes it is enough to know that transparent encryption in the Linux kernel is now implemented as a device mapper target called *dm-crypt*.

### Digression: loop devices and notes for testing

All this does not mean that the venerable *losetup* command itself is gone, or that it is useless. It is still needed in order to achieve its original function: to turn a plain file into a block device. As you read the following, and you have the urge to try things out, you can use *losetup* to experiment safely. Create a nice large file, and assign it to */dev/loop0* using the commands below.

```
dd if=/dev/zero bs=1024k count=50 of=large.file
losetup /dev/loop0 large.file
```

When you're done with all your testing, you can try the following:

```
losetup -d /dev/loop0
rm large.file
```

In all the examples below, we will be using */dev/loop0* as our block device. Once you are confident, you can substitute real devices like */dev/sda6* or whatever.

Finally, here's a quick tip: if you get strange errors when running the commands below, try running *modprobe dm-crypt* manually first.

### dm-crypt and cryptsetup

*dm-crypt* is a device mapper target whose purpose is to encrypt and decrypt data on-the-fly, as it is being written to and read from the underlying block device. Since creating and specifying device mapper tables is so cumbersome, *dm-crypt* comes with a nice command called *cryptsetup* that makes it all very easy. A typical command might look like this (you'll be asked for a password; type in anything you like but remember it):

```
cryptsetup -y create enctest /dev/loop0
```

What this does is create a mapping—you can see what it looks like by typing in *cryptsetup status enctest*. A new device called */dev/mapper/encetest* is created, and you can use it just like any other block device. You can create a filesystem on it using *mkfs* or its cousins, you can mount that filesystem somewhere, put files in it, etc. None of the data you create will be readable by accessing */dev/loop0*, or the underlying *large.file* file. When you are all done playing with this, you can remove the mapping:

```
cryptsetup remove enctest
```

The default cipher (encryption algorithm) used by *cryptsetup* is AES, the Advanced Encryption Standard, in CBC (Cipher Block Chaining) mode, with a 'plain' IV. People who are interested in what these mean are encouraged to go to Wikipedia, starting with [http://en.wikipedia.org/wiki/Disk\\_encryption\\_theory](http://en.wikipedia.org/wiki/Disk_encryption_theory), for a good introduction to all this. For the purposes of this discussion, suffice it to say that while AES and CBC are fine, a 'plain IV' is not good enough, so typically you will give *cryptsetup* the option *-c aes-cbc-essiv:sha256* to achieve better security.

There's one more aspect of *dm-crypt* that I need to point out. In order to understand this, create a mapping (make a note of the password you typed in), make a filesystem on the new device, check that it is indeed a filesystem using the *file* command, and then remove the mapping, as follows:

```
cryptsetup -y create enctest /dev/loop0
mkfs /dev/mapper/encetest
file -s /dev/mapper/encetest
cryptsetup remove enctest
```

Having done that, try the following commands a couple of times; the first time give the same password you used above, and then use some other password:

```
cryptsetup create enctest /dev/loop0
file -s /dev/mapper/encetest
cryptsetup remove enctest
```

It is reasonable to expect that the data that you wrote to the underlying device (using *mkfs*) is readable only when you type in the correct password. But you might be surprised to find that when you gave the wrong password, it doesn't report an error of any kind, and quietly returns garbage instead of your data, as seen by the result of the *file* command in each case.

This might sound like a bug, but it isn't. Technically, 'decryption' means just applying a function to the given inputs to produce an output. It cannot determine if the key is correct or not—that can only happen if there is some other way of checking the result of the decryption, perhaps

by using a checksum, or looking for a known pattern of data somewhere, etc. Naturally, the same thing happens if, in the second round, you specify some other cipher, say *-c twofish-cbc-plain*, or some other key size, like *-s 128*.

## LUKS

The fact that you have to remember what cipher algorithm you used, what its key size was, and some other details, is the biggest weakness in *dm-crypt* and *cryptsetup* from a usability point of view. And since there's no feedback when you use the wrong password or options, accessing an encrypted device you may have created months ago (like a USB disk used for occasional back-ups) becomes a bit of a game.

Ideally, you want to plug in an encrypted device, and have the system prompt you only for a password, figuring out all the rest of the information for itself. And that is precisely what LUKS (Linux Unified Key Set-up) does.

LUKS is a standard for the on-disk format of encrypted data. In actual usage, all the LUKS stuff has been rolled into the same *cryptsetup* command you saw above, with LUKS-specific options starting with *luks*. The manual for *cryptsetup* is very clear and worth a quick look.

The first step to using LUKS is to run *cryptsetup* with the *luksFormat* option on an empty device; this initialises the device with a LUKS header that contains a bunch of useful information (we'll see later what this is), and asks you to set a password to access the device. A reasonable analogy, if you don't take it too far, is that of a partition table on a disk.

Note the difference between *cryptsetup luksFormat* and *cryptsetup create*. *luksFormat* is actually writing something to the device, while *create* is doing nothing except getting ready to do on-the-fly encryption/decryption of your data—until you write to */dev/mapper/enctest*, it has no permanent effect on the disk.

Which means, once you've '*luksFormat* -ed' a device, you shouldn't do it again—this destroys the data on the disk, unlike running *cryptsetup create* with a wrong password.

So here's the canonical sequence of commands to use LUKS:

```
# one time only
cryptsetup luksFormat /dev/loop0

# each time you use it
cryptsetup luksOpen /dev/loop0 enctest
# ...work with /dev/mapper/enctest
# mkfs, mount, whatever...
cryptsetup remove enctest
```

## The LUKS header

So what does the LUKS header store?

If you try the 'wrong password' test above, you will

notice that *luksOpen* will give you three chances (this is configurable) to give the right password, but refuses to accept the wrong password. It can do this because one of the pieces of information it stored in the LUKS header when you did the *luksFormat*, is a hash of the master key to the disk, so it can easily tell whether you gave the right password or not.

You'll also notice, if you try, that *luksOpen* ignores attempts to set the cipher, key-size, etc, on the command line, because it picks up those settings from the device header.

Just out of curiosity, try the following two commands:

```
file -s /dev/loop0
cryptsetup luksDump /dev/loop0
```

The first command shows you that the system can detect an encrypted device with a LUKS header, which makes it possible to prompt for a password and auto-mount when such a device is plugged in. In fact, on most recent distributions, this is exactly what happens, so you can safely carry around an encrypted USB stick and be assured that you only have to remember the password to access it on any machine.

## Changing passwords

The second command shows you what is actually stored in the LUKS header, and you might notice eight 'key slots' in that report, with only the first one filled.

Plain *cryptsetup create* just uses the password as it is, which means that if you want to change the password, the entire disk has to be somehow re-encrypted with the new one—a complex and time-consuming process.

What *luksFormat* does instead, is to use a very long, random, value as the key to the data being encrypted. This 'master key' is then stored in the LUKS header, encrypted with your password, which is what you see in 'key slot 0'. LUKS allows you to add keys to any empty slot, and delete keys from any slot by giving the password for some other slot.

So to change the password, you simply use *luksAddKey* for the new one, and *luksDelKey* for the old one.

And why would you want multiple passwords for the same data? Well, apart from the obvious reason of allowing multiple people to use the device, here's one scenario that's worth considering. Let's say you're encrypting */home* on your desktop. LUKS allows you to use an actual file for the password, which is a pretty neat feature, if you think about it. So you create a random file on your USB stick and use it as the 'password' when doing the initial *luksFormat*. From now on, you cannot mount */home* without having that USB stick mounted (or at least, having that file accessible somewhere), so if you shut down your computer, and walk away with the stick, your data is safe.

But you now have a problem: what if the USB stick gets damaged or lost? You need some sort of insurance, and

this is what having multiple password slots gives you. Pick a nice, really long, phrase that cannot be brute-forced or guessed—something too long for everyday use. Add a key with that pass-phrase into another slot, and you have a means to get in even if the USB stick is gone. This pass-phrase should be easy to remember, or you could write it down and put it under lock and key somewhere—how you secure *that* is up to you.

This also provides protection against lost passwords. In the example above, you lost the USB stick. Before someone finds it and uses it to get access to your machine, you can use the back-up key (that long pass-phrase) to delete the other key by specifying its slot number. (It's not difficult to figure out which slot is used by which key, in case you forgot the sequence in which you created them. Every time you use a key successfully to *luksOpen*, you get a message saying something like “Key slot 2 unlocked”, so you know.)

LUKS is not just better for usability though. It also has some useful security enhancements, such as using PBKDF2 (Password Based Key Derivation Function) on the user-supplied password before using it to encrypt the master key, the ability to customise the number of iterations in PBKDF2, and the ability to spread the encrypted copy of the master key across many sectors to prevent forensic recovery of deleted keys.

## And now for something completely different!

All that said and done, this isn't what I use for most of my data now. If I were really paranoid, I would probably use my distro-supported method to encrypt the whole disk. I know what is happening under the hood anyway, in case I ever need to manually recover the data.

However, I only want to encrypt a part of my home directory (in my case, my thunderbird/spicebird mail folders), and I don't want to plan how much space they will need in advance. And resizing file systems is a pain.

What I do instead is to use EncFS. This one provides an encrypted filesystem in userspace, using the standard FUSE (Filesystem in USERspace) module. This is also called a pass-through filesystem—no pre-allocation of space is needed because the actual files are stored on the real filesystem underneath. It's great for ad-hoc encryption; for instance, when you suddenly realise you need to encrypt some files on your USB stick, there's no need to plan in advance.

It also fits into most back-up strategies; though the back-up software cannot decrypt the files, it will know which files have changed. On the downside, anyone can see how many directories and files you have, their sizes, the approximate length of each file/directory name, and ownership modes and timestamps.

If all those restrictions are a problem, this is not the method for you. However, it does a fine job of keeping my data safe from the risks that I am most concerned with, and I suspect that for most people, who are not even encrypting anything right now, this would be an excellent place to start.

Using it is pretty simple; just remember that every encrypted directory has a ‘real’ directory behind it. To start, create two empty directories and use *encfs* to mount one of them onto the other, as follows (the *encfs* command behaves differently the first time; if you read the prompts carefully it's easy enough):

```
mkdir .enc.enctest encetest
encfs $PWD/.enc.enctest $PWD/encetest
# note: the encfs command needs absolute pathnames
```

Now, let's say you copy files into *encetest*, and the contents look like this when you're done:

```
-rw-r--r-- 1 sitaram sitaram 64 2008-09-07 14:59 ./abc.txt
drwxr-xr-x 2 sitaram sitaram 4096 2008-09-07 15:00 ./cde
-rw-r--r-- 1 sitaram sitaram 183 2008-09-07 15:00 ./cde/12345.txt
-rw-r--r-- 1 sitaram sitaram 122 2008-09-07 15:00 ./def.txt
```

If you examine the real files, which are in the *.enc.enctest* directory, they might look like this:

```
-rw-r----- 1 sitaram sitaram 224 2008-06-18 00:30 .encfs5
-rw-r--r-- 1 sitaram sitaram 72 2008-09-07 14:59
./,fQHoblUNOE,1
-rw-r--r-- 1 sitaram sitaram 130 2008-09-07 15:00 ./eZj6TDV14cGxg,
drwxr-xr-x 2 sitaram sitaram 4096 2008-09-07 15:00 ./gv3VcK3nSu70J0
-rw-r--r-- 1 sitaram sitaram 191 2008-09-07 15:00 ./gv3VcK3nSu70J0/uqJ9sUjc5V35cM1tbTq8VIns
```

The most important file in this directory is the control file, *.encfs5*. If you lose that, you cannot access any of the files again. The other files are independent of each other.

As you can also see, anyone can guess the length of each of your files if they can see the encrypted files: they just have to subtract 8 from the length of each encrypted file! Plus, they can see the directory structure clearly, although they cannot guess the actual directory/file names.

As I said, these do represent compromises from a deniability point of view, as well as open up some attacks relating to information leakage. But for the simple use case of protecting the data from theft, or lending the computer to someone to use for some time, or perhaps sending it for repair, etc, this is quite sufficient.

## Last word

As the man page for *cryptsetup* says, “Mathematics can't be bribed. Make sure you keep your passwords safe!” 

**By: Sitaram Chamarty** (sitaramc@gmail.com) discovered Linux (and simultaneously, Perl) 13 years ago, and if it hadn't been for this lucky encounter he would have probably transitioned from nine years of mainframe work straight into ‘management’! Sitaram works at TCS, in Hyderabad.

# INDUSTRY NEWS

## Developers prefer commercial Linux

According to the latest Open Source Software and Linux survey by Evans Data Corporation, open source developers overwhelmingly say they would recommend using commercial versions of Linux over non-commercial versions when it comes to leading edge or high performance computing. A similar number would also recommend commercial versions over non-commercial for mission-critical large enterprise development and data centre development. Non-commercial Linux is preferred for Web development and embedded systems.

"While the open source nature of Linux and the availability of Linux source code appeal to developers doing complex leading-edge and mission-critical work," said John Andrews, president and CEO, Evans Data, "they also like the additional support they get with commercial products."

Around 400 developers active in open source development participated in the survey, which was conducted in August 2008. Here are some highlights from the survey: Ubuntu was used by 24 per cent of the open source developers and was the most used Linux distribution, although Red Hat Linux (21 per cent) and Red Hat Enterprise (19 per cent) together exceed Ubuntu's use; VMware was the top virtual machine technology used by the developers, with over a third reporting its use; Apache/BSD style licensing or GPL2 are by far the most used open source licensing models, with GPL3 and LGPL lagging far behind.

## SFLC publishes guide to GPL compliance

The Software Freedom Law Centre (SFLC), a provider of legal services to protect and advance free and open source software, has published a guide on the effective compliance with the GPL and related licences. The guide provides a basic legal overview of GPL compliance and recommends procedures that companies can implement to avoid violations. It educates the users and community of commercial distributors, redistributors, and resellers on how to avoid violations and to respond adequately and appropriately when a violation occurs.

SFLC's Bradley M. Kuhn, who has conducted GPL enforcement since 1998 and co-authored the guide, added, "Cooperative and non-confrontational enforcement has always been and remains the norm and preference of everyone in the community. Through this guide, we further advance that goal by providing even more information to help those who commercialise FOSS to comply easily from the start." The guide, entitled *A Practical Guide to GPL Compliance*, is available at [www.softwarefreedom.org/resources/2008/compliance-guide.html](http://www.softwarefreedom.org/resources/2008/compliance-guide.html)

## Red Hat, Intel plan to set up FOSS testing lab in NY

Red Hat has announced a collaboration with Intel and the City University of New York (CUNY) to form the New York City Open Source Solutions Lab at the CUNY Institute for Software Design and Development.

The Open Source Solutions Lab houses hardware provided by Intel Corporation and Red Hat Enterprise Linux software donated by Red Hat. In addition to hardware and software, the companies are also supplying support services. The lab is designed to help New York City area public sector government IT professionals consider more cost-effective and flexible technology options prior to deployment.



## Ubuntu-powered Dell PCs finally debut in India

Dell has introduced two new Vostro laptops and two desktops designed specifically to meet the needs of small businesses, governments and educational institutions operating on limited budgets in India and the world's emerging economies. The products include Dell Vostro A860, Dell Vostro A840 laptops, Dell Vostro A180 and Dell Vostro A100 desktops. According to the company, additional Vostro products designed for India and emerging economies will be introduced in the coming months.

The Dell Vostro A860 is a 15.6-inch (39.6 cm) laptop that offers an HD widescreen LCD with anti-glare coating. The device also offers the choice of a variety of Intel processors from Celeron to Core 2 Duo, and the choice of Ubuntu Linux or Windows Vista.

## Amanda wins award for best network back-up software

Zmanda, an open source back-up and recovery software provider, announced that InfoWorld has recognised its Amanda enterprise solution as the *Best Open Source Software* in the storage category. InfoWorld's annual awards, commonly known as the BOSSIES (Best of Open Source Software Awards), recognise the best free and open source software the world has to offer to businesses and IT professionals.



Amanda has been chosen as the best open source network back-up software for its extensive support of heterogeneous platforms, layered security and its exclusive use of open data formats. It is probably the world's most popular open source back-up and archiving software, protecting more than half a million servers and desktops.

# INDUSTRY NEWS



## C-DAC inks MoU with National Informatics Centre

C-DAC (Centre for Development of Advanced Computing) has signed a memorandum of understanding with National Informatics Centre (NIC), to implement BOSS Linux on select e-governance projects developed and maintained by NIC. The scope of this strategic alliance between C-DAC and NIC is to implement the latest 3.0 version of BOSS Linux, developed by NRCFOSS, in these e-governance projects.



The MoU would encourage the procurement of Linux compatible devices for use in e-governance applications. NIC will identify suitable e-governance applications for deployment of BOSS and utilise C-DAC's expertise in Linux software. C-DAC, in turn, would support telephonic, e-mail and portal delivery mechanisms for the e-governance programmes of NIC.

## Wind River acquires MIZI Research

Wind River has signed an agreement to acquire MIZI Research, a privately-held company based in South Korea that focuses on the development of mobile application platforms based on embedded Linux. With this acquisition, Wind River expects to gain access to mobile expertise that can be leveraged across its various mobile alliances as well as accelerate its mobile services presence in the Asia Pacific region. Also, a number of fast-growing consumer electronic market segments are turning towards Linux, and Wind River aims to pitch itself to this growth in the mobile Internet devices, automotive infotainment and mobile handsets industries.

Wind River will pay up to \$16 million in cash to acquire substantially all of the outstanding shares of MIZI. The completion of the acquisition is expected to occur in Wind River's third fiscal quarter ending October 31, 2008.

"Wind River is targeting the mobile device market as a strategic growth opportunity within our Linux business. With the addition of the MIZI team to Wind River, we will benefit from their mobile expertise in areas such as telephony, feature-rich user interfaces and multimedia, as well as their world-class mobile systems integrator expertise, complementing our current efforts across a variety of mobile alliances," said Ken Klein, chief executive officer, Wind River.

## Linux Foundation's first End User Collaboration Summit

The Linux Foundation has announced the first Linux Foundation End User Collaboration Summit. The inaugural summit will take place between October 13 and 14, 2008, in New York and will provide end users a direct connection and voice to the kernel community. The summit will offer an opportunity for end users to learn and interact with leaders from within the Linux community, including the highest level maintainers and developers.

The inaugural summit will also give Linux community maintainers and developers direct access to knowledge sharing opportunities with the end users. The event will include Q&A sessions, one-on-one discussions, and presentations by some of the important names in the Linux community like Novell CEO Ron Hovsepian, and Paul Cormier, executive vice president and president, products and technologies, Red Hat.

## GNU turns 25

The GNU operating system turned 25 this month and the Free Software Foundation (FSF) had kicked off a month-long celebration for the anniversary by releasing 'Happy Birthday to GNU', a short film featuring the English humorist, actor, novelist and filmmaker, Stephen Fry.

In the five-minute film, Fry compares the free software operating system to 'good science' and contrasts it with the 'kind of tyranny' imposed by the proprietary software produced by companies like Microsoft and Apple. He encourages people to use free GNU/Linux distributions like gNewSense and free software generally, for freedom's sake.

Peter Brown, the FSF's executive director, said on the occasion, "We intend for the 25th anniversary to be more than just a reflection on the history of the free software movement, because despite all of the success brought about by the GNU system and other free software projects, we still need a determined effort to replace or eliminate the proprietary applications, platforms, drivers and firmware that many users still run. In this light, the video of Stephen Fry is not just a celebration, but a rallying call for the work that still needs to be done."

It was in the month of September 1983 that Richard Stallman introduced his plan of releasing free software based on a UNIX-like system, available with a free open source code, allowing anyone to modify and redistribute it. Today, over 300 software packages are released under the GNU Project, and new programs are being added all the time. The video 'Happy Birthday to GNU', along with more information about GNU software and its philosophy, is available at <http://www.gnu.org>





# INDUSTRY NEWS

## Red Hat buys Qumranet to extend its virtualisation market

Red Hat, in its effort to transform the virtualisation market, has acquired a privately held open source software company, Qumranet, for \$107 million. The acquisition includes Qumranet's virtualisation solutions, its KVM platform and SolidICE offering, a virtual desktop infrastructure (VDI), which together present a comprehensive virtualisation platform for enterprise customers. In addition, Qumranet's team of professionals that develop, test and support Qumranet solutions, and its leaders of the open source community KVM project, will join Red Hat.

The Qumranet acquisition also extends Red Hat's virtualisation solutions for managing Windows desktops. SolidICE is a high-performance, scalable desktop virtualisation solution built specifically for the virtual desktop. SolidICE is designed to enable a user's Windows or Linux desktop to run in a virtual machine that is hosted on a central server. It is based on the industry-leading Simple Protocol for Independent Computing Environments (SPICE) protocol, which overcomes key barriers to VDI adoption, including a superior user experience enabled by the SPICE protocol capabilities.

The acquisition will help Red Hat to drive comprehensive virtualisation technology and management solutions into every system, from servers to desktops, on both Linux and Windows. Red Hat's solution components include an embedded hypervisor, which supports all major operating systems, a management platform for both virtual and physical systems, cloud and grid management solutions, clustering solutions, and integrated security infrastructure.

## Cisco to acquire Jabber

Cisco has announced its intention to acquire privately held Jabber, Inc, a provider of presence and messaging software. Based in Denver, Jabber will work with Cisco to enhance the existing presence and messaging functions of Cisco's Collaboration portfolio.

The acquisition will enable Cisco to embed presence and messaging services 'in the network' and provide rich aggregation capabilities to users through both on-premise and on-demand solutions, across multiple platforms including Cisco WebEx Connect and Cisco Unified Communications.

Jabber provides a carrier-grade, best-in-class presence and messaging platform. Jabber's technology leverages open standards to provide a highly scalable architecture that supports the aggregation of presence information across different devices, users and applications. The technology also enables collaboration across many different presence systems such as Microsoft Office Communications Server, IBM Sametime, AOL AIM, Google and Yahoo. Jabber's platform leads the market in system robustness, scalability, extensibility and global distribution.



## Microsoft, Novell offer joint virtualisation solution

Microsoft and Novell have announced the availability of a joint virtualisation solution optimised for customers running mixed-source environments. The joint offering includes SUSE Linux Enterprise Server, configured and tested as an optimised guest OS running on Windows Server 2008 Hyper-V, and is fully supported by both companies' channel partners, including Dell. The offering claims to provide customers with the first complete, fully supported and optimised virtualisation solution on both Windows and Linux environments.

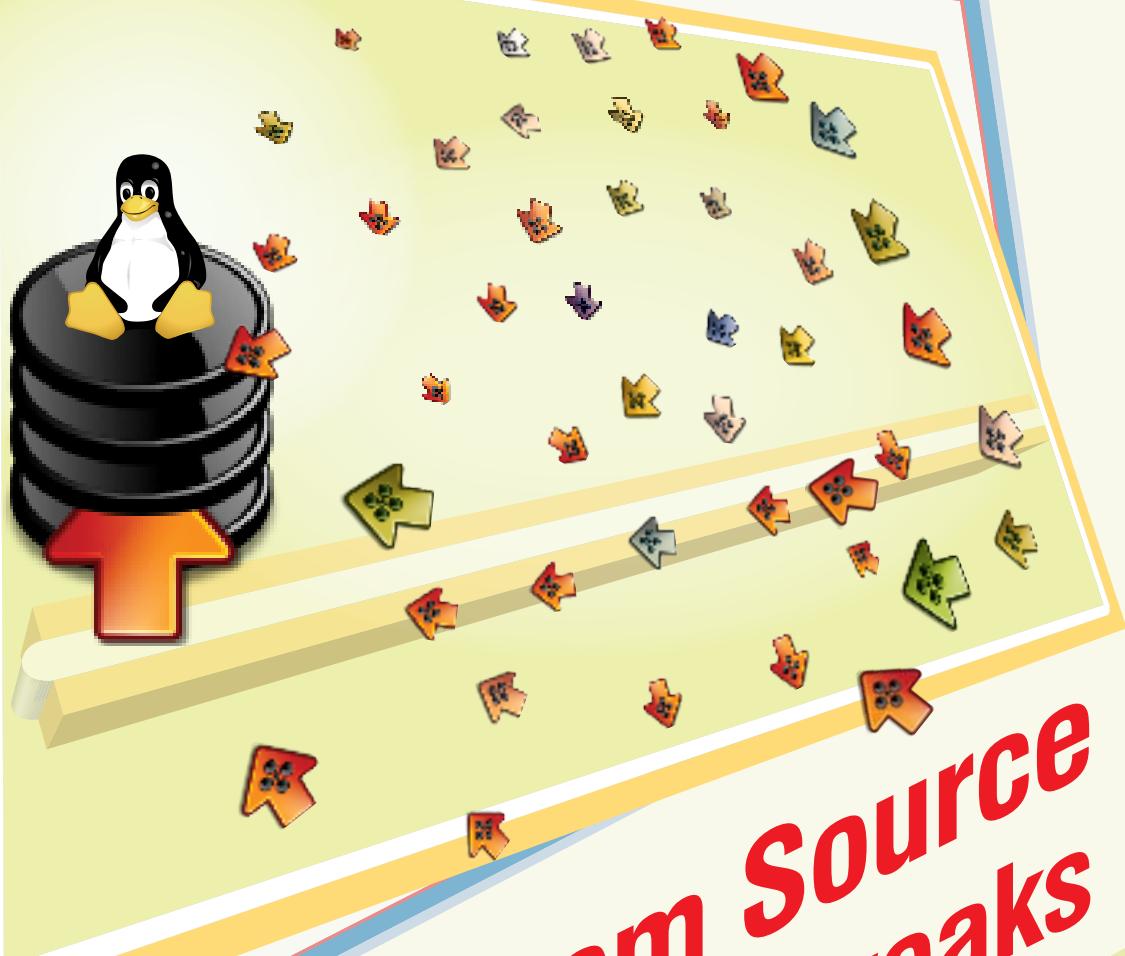
The virtualisation solution is the first to include technology developed by both companies at their joint Interoperability Lab, including virtual machine adapters built to optimise SUSE Linux Enterprise Server as an optimised (or as it is often referred to as 'enlightened') guest operating system on Windows Server 2008 Hyper-V, providing optimised performance to SUSE Linux guests.

## Ingres database to help LYNX services gain quick claims



LYNX Services, an automobile glass insurance claims management company, has adopted the open source version of Ingres Database to process auto glass claims for their clients. Ingres Corporation provides high availability, quick transactional processing and an easy deployment language that helps LYNX Services manage more than three million claims annually. More than 6,000 auto glass retailers depend on LYNX Services for electronic data interchange transactions.

## Optimising Performance on Linux Desktops —Part 3



# Compiling From Source and Kernel Related Tweaks

Here are various advantages, disadvantages and tips on how we can go about compiling applications from source, including the kernel.

In our quest to achieve the maximum performance, so far we have already taken a look at doing installations the right way (in Part 1, published in the July 2008 issue of *LFY*) and moved on to Fluxbox (in Part 2, published in *LFY*'s August 2008 issue). This article intends to direct every desktop user looking for more performance towards the advantages of compiling programs and the kernel from source.

Why compile programs from source?

- When a program is compiled from source it can be optimised for the hardware it will be running on. This would make the program deliver maximum performance.
- It can also be customised with regard to the other software installed on the system.
- Unwanted/unused components from the software can be removed and additional components can be installed, thus making it lighter or more feature-enriched.
- Upgrading to newer releases of programs can be done instantly by just downloading the latest source version and compiling it, rather than waiting to download it from the distribution's software repositories.
- There is nothing that is unavailable. Every FOSS application built has its source available. Never would you have to complain about packages missing from your distribution's repository.

However, as it often happens when there are advantages, there are disadvantages too. The following are a few:

- Source packages take a very long time to compile. This can be a big issue if you are dealing with a package that comes up with regular updates or fixes.
- The hardware requirements and compile times are inversely proportional. That means the better the hardware, the faster you would be able to compile a program.
- Dependency issue fixing is a little more tedious, unlike using a distribution with amazing package management skills like Debian.
- Imaging an OS for multiple machines with different hardware is not possible since each OS will have to be tailored for that particular hardware.
- Disk space used is more than that by pre-compiled binaries.

I'm sure that I am missing a few advantages as well as disadvantages, but this is basically a fundamental overview of what to expect, and what not to, from source compiled programs.

## Taking the easy way out

There are distributions like Gentoo and LFS that compile every single package from source and allow the user to customise the OS as required. The greatest advantage derived from these distributions is the ability to tailor

## Why a Debian-based distro?

There are primarily two reasons for choosing a Debian-based distribution:

- The vast Debian repositories put users in the favourable position of rarely having to compile from source—if you have a low-end machine this helps a lot!
- It offers simple ways to create installable packages from the compiled sources—a file with a .deb extension.

every bit of it from scratch. However, since compiling requires good hardware and lots of time, I would not opt for one of these distributions on a low-end machine. Moreover, many of the applications, even after compiling from source, do not make much difference to the overall performance of the machine. For example, for a package like *ssh* (the *openssh* server, that runs as a daemon), it does not make a wee bit of difference even after it is compiled from source.

Of course, you can install and run applications compiled from source on distributions that install pre-compiled executables from packages. Using such a distribution is the best way to ensure that not too much time is spent on compiling programs that won't be beneficial in improving overall performance. In this way, desktop users can be very selective in choosing which packages they would want compiled from source, and which ones they would install from the available pre-compiled package repositories. This enables the user to enjoy the benefits of both worlds!

Distributions along with pre-compiled binaries also provide source packages of it. This is because major distributions do a little bit of modification to many of the original source packages in order to maintain compatibility and integrate properly with the other packages of the distribution.

Taking care of these differences while building source packages that are not obtained from the distribution can be a time consuming job. Thus, the easiest way to compile source packages for performance is to get them from the distributions repository itself and then build them as per requirement.

## Packaging

Having installable binary packages has its advantages:

- Packages can be created on one machine and installed on another, easily.
- Installing as well as uninstalling packages is simple.
- Upgrading from an old to a newer version of a package is a breeze.

Source code usually comes in tarballs and packages made for Debian-based distributions are called debs. These debs handle the job of placing the required files in the right locations, running scripts, restarting the required services, etc. Packages are categorised as per processor architecture.



kernel 2.6.23 onwards. Since v2.6.24, the default memory allocator (SLUB) and the default scheduler, along with various other improvements, help in providing the desktop user with optimum performance.

I haven't noticed any significant difference in performance with a recompiled kernel for the i686 architecture, which implies recompiling a kernel for performance for such processors is a waste of time. The same experiment when carried out on an older processor gave different results: the system's responsiveness improved significantly with a recompiled 2.6.24 kernel on an i586.

To decide whether recompiling could improve your desktop performance, you will have to, of course, go ahead and recompile one. Given below are the steps to do so, and also a few options that could make you feel the difference. There are two recommended locations to download sources from—either from *kernel.org*, or from the distribution's repositories. I would recommend using the source obtained from the distributions repositories because each distribution adds or removes a few functionalities to/from the stock kernel. In case you are using one of the exclusive distribution-added functionalities, and recompiled stock kernel does not include function, *that* would result in the system not performing in the desired manner.

The reasons a desktop user would want to recompile and use a stock kernel on a desktop are:

- For the latest drivers
- For the most recent security fixes in case your current distribution isn't updating it.
- The kernel source that the distribution provides is not 2.6.24 (or has not back-ported functionalities from recent kernels).
- Simply for testing purposes.

Before delving any deeper into kernel compilation, I would recommend you to go through the well-written book by Greg Kroah-Hartman called "*Linux Kernel in a Nutshell*", available under the Creative Commons Attribution-ShareAlike 2.5 license at [www.kroah.com/lkn](http://www.kroah.com/lkn).

## Retrieving the kernel sources

The stock kernel sources can be downloaded from *kernel.org*. At the time of writing, the latest stable version of the Linux kernel was 2.6.26.2.

```
$ wget http://kernel.org/pub/linux/kernel/v2.6/linux-2.6.26.2.tar.bz2
$ tar -jxf linux-2.6.26.2.tar.bz2
$ cd linux-2.6.26.2
```

From now on *linux-2.6.26.2* will have to be considered as the working directory. In case you want to use the distribution's kernel sources, first find which kernel is being used, by issuing:

```
$ uname -r
```



Figure 5: The menuconfig screen. The top most section of this screen gives information on how to navigate through the configuration options and other important functionalities.

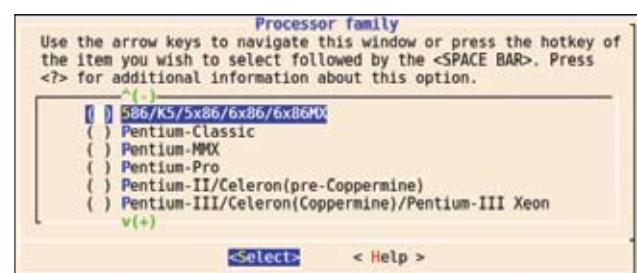


Figure 6: A very important option to select is the processor for which the kernel is being built. Located in Processor type and features→Processor family.

This will output a line with the kernel version. In my case it is 2.6.24-20. The revision number in this case is '20' which is information we would not require.

Then install the source and extract it from the archive:

```
$ sudo apt-get install linux-source-2.6.24
$ cp /usr/src/linux-source-2.6.24.tar.bz2 .
$ tar -jxf linux-source-2.6.24.tar.bz2
$ cd linux-source-2.6.24
```

Remember, we intend to package the compiled kernel and make it portable; hence, it is not required to compile the kernel on the same machine you plan to install it on. Preferably, use the most powerful machine you can get your hands on with a latest copy of a Debian-based distribution installed on it. Because different GCC versions produce code of different quality and compactness, choosing the most recent one is always the safer bet.

The next step is to install all the programs that are required to compile the kernel:

```
$ sudo apt-get install kernel-package libncurses5-dev fakeroot
```

## The actual compilation

There are various methods to compile a kernel; many of the steps used are common in all these methods. Here, I will

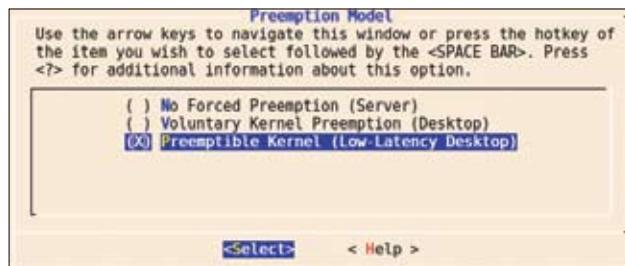


Figure 7: The most important option, according to me, for a desktop user. Located in Processor type and features→Preemption Model. This is very useful for multimedia applications. It compromises a little on the performance front, but definitely increases responsiveness.

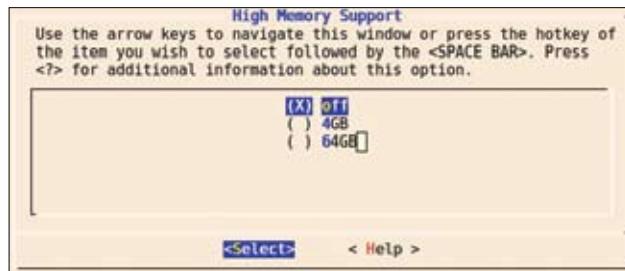


Figure 8: An example of things that can be disabled. If your target system does not have more than 4GB of RAM, this option can be safely turned off. Located in Processor type and features→High Memory Support.

briefly describe a method that will work for both sources (the distribution as well as stock).

A few tips:

- Compile the kernel in your home directory as a normal user.
- If you are unsure of any options, choose defaults.
- Know for what requirements you are compiling the kernel, especially hardware.

Configuring the kernel from scratch is a time consuming job. To make it easier, we will take the distribution existing kernel *config* file and modify it.

```
$ cp /boot/config-2.6.24-20-generic .config
```

In my case, since my kernel version is 2.6.24-20, that gets appended to the filename. Select according to the kernel you are using.

```
$ make-kpkg clean
```

```
$ fakeroot make-kpkg --initrd --revision=dg.1 \
--config=menuconfig kernel_image kernel_headers
```

The above command opens up a window that allows you to configure options. Here the *--revision* flag is used to add a number to the current version of the kernel. The above command will generate two packages—the kernel image and the header files for the same version. The headers package is of importance while compiling additional modules later. In our case we are using the *ncurses*-based configuration utility over many of the

impressive graphical ones, because it's easy to use with simple keystrokes, and too many additional packages need not have to be installed.

There are three choices for every option in a section: to include the code built into the kernel, to include the code as a module to the kernel, or to not include the code at all.

Most of the code is compiled to be modular; in this way, the kernel loads the module when required and unloads it after its requirement is over. This saves a lot of RAM as compared to having everything built into the kernel, since all code is loaded to RAM.

It is noticed that every time a module is loaded into RAM the system suffers a performance hit; hence the idea is to compile as much as possible as built-ins. Just in case you feel you can obtain gains by having fewer modules, have only those modules compiled as built-ins that the system absolutely needs—for example, the filesystem drivers. In my case the */boot* partition is ext3 and the others XFS—my system will refuse to boot without having these modules loaded. Likewise, these drivers would be good candidates to be built into the kernel.

In case the RAM on the system is considerably less, having the kernel compiled with options like virtualisation is of no use. Hardware like graphics cards and sound cards are rarely changed, thus choosing the right one and disabling the others is very much recommended. It is very important to know existing hardware and related driver names. Options like SELinux and AppArmor located in 'Security Options' if not implemented, can be disabled. The less options you enable, the shorter the time it will take to compile the kernel.

After configuring the kernel, save the configuration and exit the configuration utility. The compilation will carry on from then on. Once the entire process is completed, two files will be created in the directory above the working directory. These will have to be installed.

If the kernel is supposed to be used on the same system on which you compiled it, issue the following command to install it:

```
$ sudo dpkg -i ../linux-*.deb
```

Else move the packages to the appropriate system and install them as the root user.

That's it! Reboot and try out the newly compiled kernel. If this was your first try at recompiling a kernel and if you see that things are not working as desired, then don't be disheartened—there are a lot of resources available that can help you solve almost any problem related to kernel compilations.

You may not get a perfect kernel the first time, but it's important to know that if you stop trying, you may never will. Do not uninstall the older kernels that had been installed by the distribution—they serve as back-ups to boot your system with when your compiled kernel fails.

However, if your customised kernel doesn't work, you don't need to keep that too—uninstall it by booting into the older kernel and issuing the following command:

```
$ sudo dpkg --purge linux-image-“version number” \
linux-headers-“version number”
```

Here, the 'version number' will be of the newly installed kernel. After uninstallation, delete the created packages, restart from copying the configuration file in case you intend to start from scratch again; or from the step after that, in case you only need to do changes in the configuration file.

After many recompiles and extensive usage you may feel that just by choosing the pre-emption model as low latency, the responsiveness is sufficient and all the other optimisation options provide very little or no difference in performance—thereafter, you can continue using a pre-compiled kernel from the Ubuntu repositories. The Ubuntu repositories have a pre-compiled kernel called 'linux-rt' that can be installed and used instead of the generic one. To do so, issue:

```
$ sudo apt-get install linux-rt
```

The greatest disadvantage of having a custom kernel is that, to obtain the latest security updates and features,

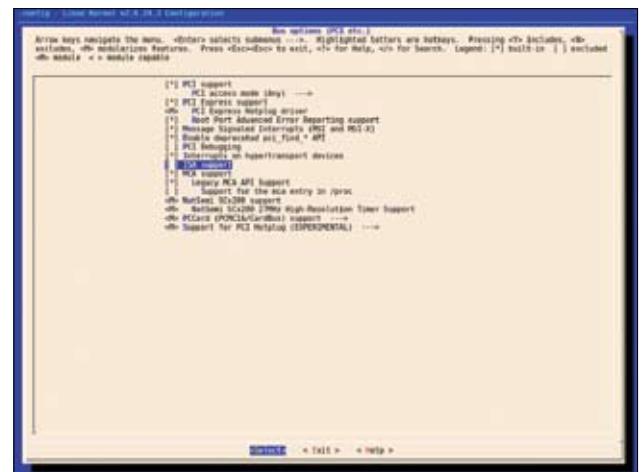


Figure 9: Another example of an option that can be disabled if the target system does not have support for ISA cards.

a newer version of the kernel will have to be manually recompiled every time. In that case, the only systems to have recompiled kernels are those that have an unsupported OS (because of having gone past the support period).

## Patching a kernel

As I mentioned earlier in the article, at times there may be problems with the current kernel and you may find code that does a better job or adds a new feature. To use this code you will need to patch the existing kernel

# 2<sup>nd</sup> Anniversary

# IT Infrastructure Services  
# Web Presence Services  
# Technical Training  
# SysAdmin Staffing

 **The UX Group**  
A SysAdmin Company

# 40-1-92, Khanna Nagar  
Near Benz Circle, VIJAYAWADA - 10  
Ph: +91-866-6662368, 2497499  
Mob: +91-9989777150

<http://theuxgroup.com>

## SEEKING LINUX GEEKS



Rhythm & Hues is an Academy Award® winning animation and VFX studio that works on major Hollywood feature films including "The Incredible Hulk" "The Chronicles of Narnia" and many more.

We are recruiting for various roles in our technology team at our Mumbai and Hyderabad facilities.

Candidates should have some Linux exposure. Positions open for Freshers too.

For more details visit: [www.rhythm.in](http://www.rhythm.in)

Mail your resume to: [techjobs-ind@rhythm.com](mailto:techjobs-ind@rhythm.com)

**Rhythm & Hues**  
STUDIOS INDIA PRIVATE LIMITED

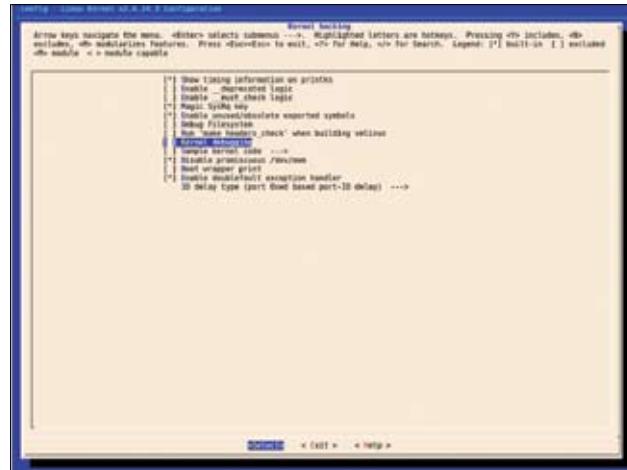


Figure 10: This probably is the most important option to disable if you are not planning to do any kernel related development. Enabling this option increases the size of the kernel. On a desktop machine, any option that includes 'debug' in its name can be safely disabled. Located in Kernel Hacking → Kernel debugging.

source to include it. Patching is also another way to update the kernel sources to a higher version. Rather than downloading a whole kernel source, you could just download a tiny patch and apply it to the current source to update it.

In the example below, a pre-patch 2.6.27-rc3 is applied to the previous full stable release, which is 2.6.26. This is done to update the existing 2.6.26 kernel source to alpha release 2.6.27-rc3. Pre-patches can only be applied to the previous full stable releases of the kernel. The latest full stable release at the time of writing is 2.6.26 and the latest prepatch is 2.6.27-rc3. For more information about pre-patches refer to [kernel.org/patchtypes/pre.html](http://kernel.org/patchtypes/pre.html).

Download and unpack the pre-patch into the 2.6.26 source directory.

```
$ wget http://www.kernel.org/pub/linux/kernel/v2.6/testing/
patch-2.6.27-rc3.bz2
$ bzip2 -dc patch-2.6.27-rc3.bz2 | patch -p1
```

Now, compile the kernel!

## A few more tweaks

The `/proc` filesystem is a virtual filesystem that resides in RAM and does not consume any space on the disk. Almost all information about the running system can be found in files from the `/proc` directory. Many of the kernel parameters can be manipulated in real time to increase the efficiency of the system.

Most of these values are stored in files in the `/proc/sys/` directory. Since all the contents of the `/proc` directory are stored in RAM, the changes made to these files go back to the previous state after a reboot. To get these settings working across reboots, it can be stored in the `/etc/sysctl.conf` file. The following are a few of those settings that can help in improving the efficiency of the system:

```
vm.swappiness=15
```

From [kerneltrap.org](http://kerneltrap.org): “Swappiness is a kernel ‘knob’ (located in `/proc/sys/vm/swappiness`) used to tweak how much the kernel favours swap over RAM; high swappiness means the kernel will swap out a lot, and low swappiness means the kernel will try not to use swap space.” Any value in the range of 10 - 20 can be considered good for a desktop system.

```
vm.vfs_cache_pressure=40
```

This value controls the tendency of the kernel to reclaim the memory that is used for caching of the directory and inode objects. The default value is 100; lowering the `vfs_cache_pressure` value causes the kernel to prefer to retain dentries (directory entries) rather than reclaim it. A dentry describes a name of a file: the inode plus the pathname used to find it.

For the values to get applied immediately after editing the `/etc/sysctl.conf`, the following command should be issued:

```
$ sudo sysctl -p
```

It's simple to understand and generate newer variables. To understand the nomenclature of entries in the `/etc/sysctl.conf` file and how they translate to locations in the `/proc` filesystem, here's the explanation: the dots in the variable name are converted to forward slashes in `/proc`. For example, a variable called `vm.swappiness` would mean the value stored in the corresponding `/proc/sys/vm/swappiness` file.

For more information, refer to the kernel documentation that can be downloaded from the repositories as a package:

```
$ sudo apt-get install linux-doc
```

All the documentation related to the `proc` filesystem is stored in `/usr/share/doc/linux-doc-”version number”/Documentation/filesystems/proc.txt.gz`

## That's all for now

Hope this article helps you keep your existing OS in sync with the most recent kernel always. With every new kernel comes a little increase in performance, greater security and more features – so, what are you waiting for? 

**By: Stanley Thomas.** The author is a roving Linux systems administrator with a solutions company in Goa. He has a craze for Linux desktops and their performance. He supports Linux server and desktop set-ups for clients, including mixed Linux-Windows environments, and can be contacted at [stanleythomas1@gmail.com](mailto:stanleythomas1@gmail.com).

# No one brews Coffee & Conversation Like Karan.



All the madness, all the masala, all the gossip and all the controversies now captured in special edition books.

Relive those fun moments with your favourite stars where they revealed it all. Read through those candid interviews that provoked conflicts and kindled friendships in tinsel town.

Review on People Magazine on Aug 15, 2008



“Koffee with Karan is a racy and fun read. The book is well laid out with enticing blurbs that lure you to read on. ”

Mr. Bachchan has been the biggest ever superstar in Hindi Cinema  
-Shahrukh Khan

I washed everyone's clothes including underwears  
-Rakhi Sawant

I go to the gym so that I can watch men working out!  
-Lara Dutta

Size does matter! -Malaiika Arora

I was great in bed -Saif Ali Khan



Available in  
Hindi & English

Shop online at  
[www.dpb.in](http://www.dpb.in)

STAR ★ Koffee with Karan Logo and all associated art © STAR India Pvt Ltd. All Rights Reserved.

Images used in the advertisement are from the show

# The Building Blocks of Your Drupal Website

Adding multiple features to your website in order to give it a unique identity is often a tedious task. It requires a lot of effort on account of architecture and coding, while making a significant dent in your pocket too. But with several contributed 'modules' and 'themes' available in Drupal, all this can be done in minutes.

Welcome back! Hope you enjoyed the first session in the series that was published in the August 2008 issue of *LFY*. By now you know what is required to install and configure a Drupal-based website. In this session we will take a dip into understanding the Drupal terminology and get better acquainted with the available options in Drupal.

## The Drupal terminology

First things first: Let's get started by getting ourselves accustomed to the terminology used by Drupal, quickly.

- **Node:** All content created from the *Create Content* menu item is stored in a 'node' table in Drupal's database, and is termed as a node. A node can be a story, a page, an audio clip, video, etc.
- **Modules:** A module is a piece of



code or program written in PHP that extends Drupal's features and resides in the *modules* folder after the installation. It is a kind of plug-in, generally written to provide specific features such as private message, blogging, etc. It uses the Drupal framework and follows Drupal coding standards. The default installation comes with core modules like a block, node, user, etc, that provides core features like user registration, block management and many more. Contributed modules are those that are not part of a standard install, but can be downloaded from the Drupal website and work seamlessly with the Drupal core to provide specific features and/or functionality.

- **Themes:** This is a collection of PHP, CSS and image files that determine the look and feel of the website. Theme files reside in the *themes* folder of the Drupal install. By default, Drupal comes with a PHP Template engine (a collection of PHP functions), that reads these files and converts them into HTML. Themes are also like plug-ins and can be downloaded from *drupal.org*. A website's look and feel can be changed by choosing a theme from the admin theme settings section.
- **Block:** This is a piece of content, data or navigational link that can be positioned within a page. It generally comprises the title and content. Its placement in a page is controlled via admin settings.
- **Menu:** A menu is a clickable link on a page that can be configured through admin settings. The Drupal core uses menu systems for handling all requests. Drupal's menu system determines which function to call, based on the URL or the requested page.
- **Clean URL:** URLs without "?" (question marks) are often referred to as Clean URLs. By default, when Drupal is installed, you will find links with "?" in it—for example, '?q=user/register', '?q=node/12'. You may want to remove this question mark from the URL to make it more human readable and easier to remember. In order to do that you need to run the *Clean URL Test* and enable the *Clean URLs* setting from *Administer → Site Configuration → Clean URLs*.
- **Cron:** This is a call to the *cron.php* program that resides in the root folder of your Drupal installation. It is usually set up by your website's administrator in the

Core - required				
Enabled	Name	Version	Description	
<input checked="" type="checkbox"/>	Block	6.2	Controls the boxes that are displayed around the main content.	
<input checked="" type="checkbox"/>	Filter	6.2	Handles the filtering of content in preparation for display.	
<input checked="" type="checkbox"/>	Node	6.2	Allows content to be submitted to the site and displayed on pages.	
<input checked="" type="checkbox"/>	System	6.2	Handles general site configuration for administrators.	
<input checked="" type="checkbox"/>	User	6.2	Manages the user registration and login system.	

Figure 1: Core modules listing

Core - optional				
Enabled	Name	Version	Description	
<input type="checkbox"/>	Aggregator	6.2	Aggregates syndicated content (RSS, RDF, and Atom feeds).	
<input checked="" type="checkbox"/>	Blog	6.2	Enables keeping easily and regularly updated user web pages or blogs.	
<input checked="" type="checkbox"/>	Blog API	6.2	Allows users to post content using applications that support SPML-RPC blog APIs.	
<input checked="" type="checkbox"/>	Book	6.2	Allows users to structure site pages in a hierarchy or outline.	
<input type="checkbox"/>	Color	6.2	Allows the user to change the color scheme of certain themes.	
<input type="checkbox"/>	Comment	6.2	Allows users to comment on and discuss published content. Required by: Forum (disabled), Tracker (disabled)	
<input type="checkbox"/>	Contact	6.2	Enables the use of both personal and site-wide contact forms.	
<input type="checkbox"/>	Content translation	6.2	Allows content to be translated into different languages. Depends on: Locale (disabled)	
<input type="checkbox"/>	Database logging	6.2	Logs and records system events to the database.	
<input type="checkbox"/>	Forum	6.2	Enables threaded discussions about general topics. Depends on: Taxonomy (disabled), Comment (disabled)	
<input type="checkbox"/>	Help	6.2	Manages the display of online help.	

Figure 2: Optional modules listing

*Control Panel* of your Web server to execute at specific intervals of time. It executes a command or even a script (a group of commands), often referred to as 'Cron Jobs', from all Drupal modules periodically.

- **CCK (Content Construction Kit):** A module, along with several other modules, allows website developers to define their own content types and custom fields such as date, images, videos, etc.
- **Taxonomy:** The word 'taxonony' means "the science of classification". This module provides Drupal users a very flexible tool to hierarchically categorise and tag content. A category or tag assigned to a node is referred to as 'term', and a group of terms is known as 'vocabulary'.

Let's now start from where we'd left off earlier and perform some simple tasks. The following section assumes that you have already installed a Drupal website and you are logged in as the first user.

**Enterprise Mail Server, Linux SBS Server, Anti SPAM, Antivirus and HTTP Filtering.**

**Bandwidth Management, Internet Access Control, Content Filtering, Web Access Reporting.**

**Customised Linux Product Development.**

**TechnoInfotech™**

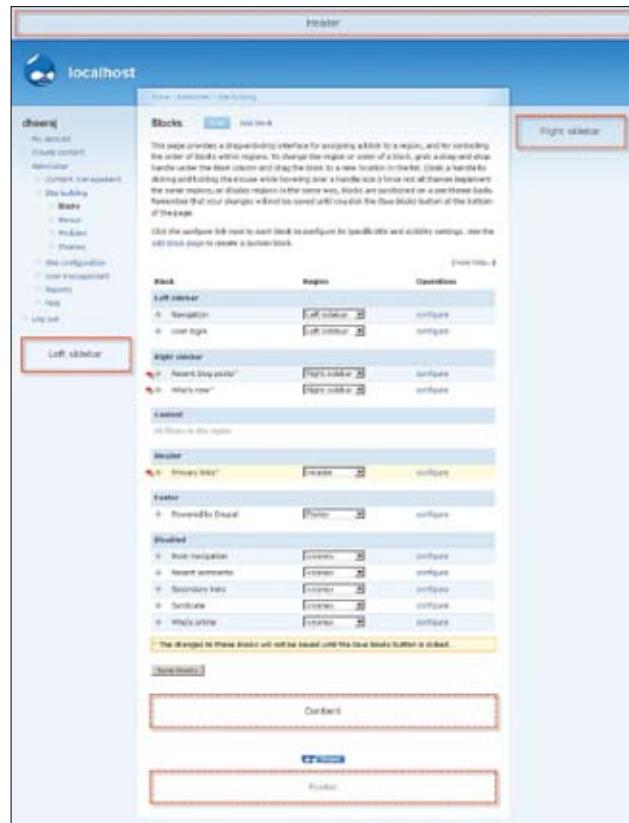


Figure 3: Blocks listing page

## How to enable a new feature? Install a module

Let's add a book and blog feature on our site. A book is a collection of hierarchically arranged Web pages, and a blog is where users can publish their thoughts.

Let's go to *Administer*→*Site Configuration*→*Modules*; you will be presented with a list of all available modules classified as Core-Required and Core-Optional with a checkbox for every module (refer to Figure 1 and 2). Parse the list and choose a checkbox next to the *Book* and *Blog* modules. Next, click the *Save Configuration* button. You will be presented with a confirmation message. The confirmation message is proof that these modules have been installed and all necessary backend/database settings have been taken care of.

Now, you can add the book and/or blog content to your website. To get started, click on *Create Content*. You will be presented with a page with two more options *Blog* and *Book*. Next, click on *Blog* and fill up the presented form. Make sure to select *Publish* and *Promote to Front Page* under *Publishing Options* before you click the *Save* button. Click on the *Home* link to see if your blog item is listed there. Isn't it great? Hey, you just managed to add a new feature to your website.

## Do it yourself

Enable the *Book* module from *Administer*→*Site Configuration*→*Modules*. Go to *Create Content* and add your book's pages. Use *Book Outline* to organise the

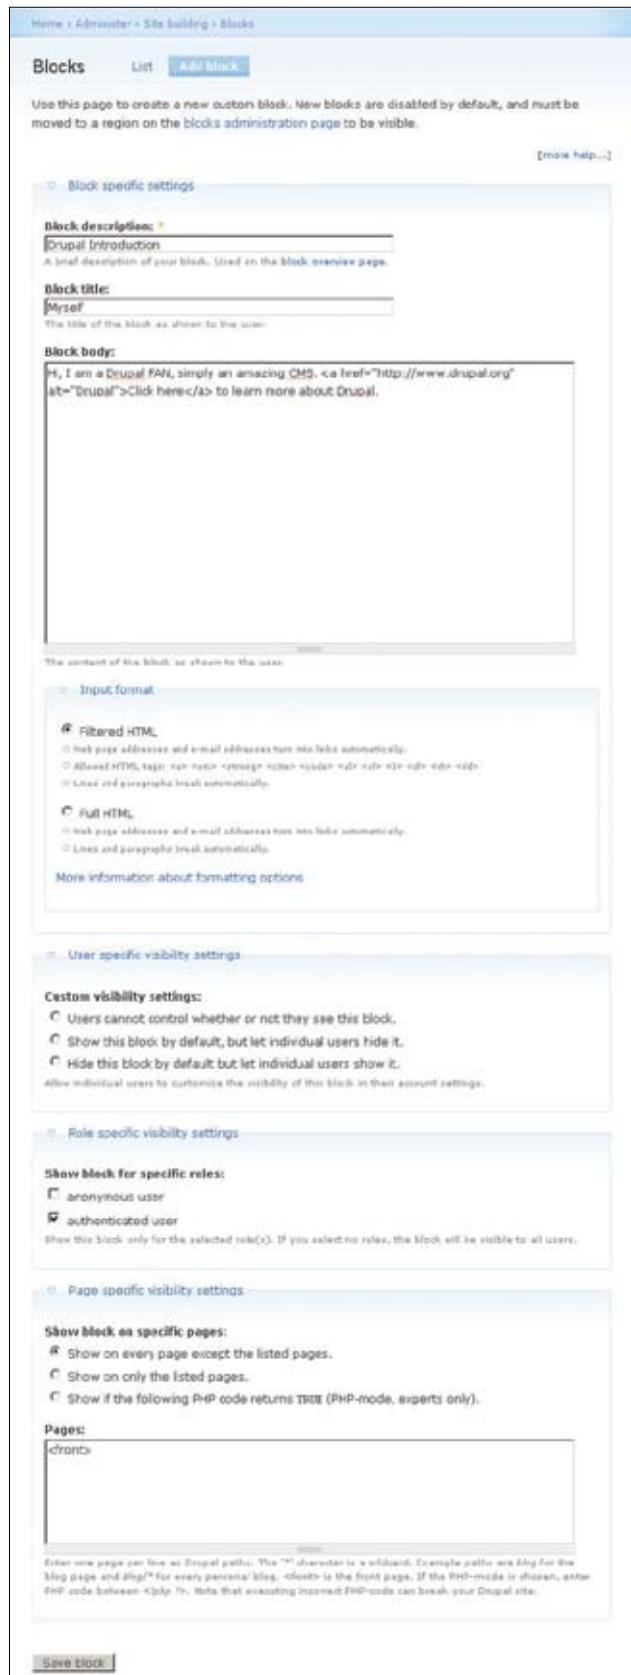


Figure 4: Add a new block

hierarchy of the pages. Similarly, follow *Create Content*→*Page* to create static pages—viz., 'About', 'Product',

Menus List menus Add menu Settings

Enter the name for your new menu. Remember to enable the newly created block in the blocks administration page.

[more help...]

**Menu name:**

The machine-readable name of this menu. This text will be used for constructing the URL of the menu overview page for this menu. This name must contain only lowercase letters, numbers, and hyphens, and must be unique.

**Title:**

**Description:**

**Save**

Figure 5: Add new menu page

Important Links [List items](#) [Add item](#) [Edit menu](#)

To rearrange menu items, grab a drag-and-drop handle under the Menu Item column and drag the items (or group of items) to a new location in the list. (Grab a handle by clicking and holding the mouse while hovering over a handle icon.) Remember that your changes will not be saved until you click the Save configuration button at the bottom of the page.

Menu item	Enabled	Expanded	Operations
<input checked="" type="checkbox"/> About Me	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<a href="#">edit</a> <a href="#">delete</a>
<input checked="" type="checkbox"/> Google Search*	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<a href="#">edit</a> <a href="#">delete</a>

\* Changes made in this table will not be saved until the form is submitted.

**Save configuration**

Figure 7: Menu item list page



The screenshot shows a Drupal site at 'localhost'. The main content area displays 'My First Blog' with a recent post. In the sidebar, there is a 'Recent blog posts' block and an 'Important Links' block. The 'Important Links' block contains links for 'About Me' and 'Google Search'. A red arrow points to the 'About Me' link in the sidebar.

Figure 8: Check out what your site will now look like

and quick to add new features? Generally, all Drupal modules come with a *README* and/or *INSTALL* file that contains descriptions of the features and instructions on how to use it—make sure you read them.

## Configuring and managing blocks

As discussed earlier, a block is a piece of content or data. Blocks can be managed in the administrative section of the site. Go to *Administer*→*Site Configuration*→*Blocks*, and you will be presented with a list of all the available default blocks and the available content regions. (Refer to Figure 3). Notice that this page also marks the position of available content regions in dotted boxes (header, left sidebar, right sidebar, content and footer). Click on the plus image next to 'Recent blog post' list item from the 'Disabled' section, and drag and drop it to the right sidebar in the list. Similarly, drag and drop the 'Who's new' list into the right sidebar section, and the 'Primary Links' list to the header section. You will notice an \* (asterisk) next to each of these items indicating that the data is not yet saved. Click on the *Save Block* button to save these settings. See the effect of these settings in your site.

Apart for the blocks that are available by default, you can also add your own blocks. On the Admin Blocks page mentioned above, you will find an *Add Block* tab. Click this, and on the page thus presented, specify Block Title, Description and Body. The 'Body' field is where you specify your content that will be displayed on your website. Let's add simple text like: "Hi, I'm a Drupal fan. Drupal is simply an amazing system." You can choose to provide simple text or specify an HTML. If you use the *filter HTML* option (from the input format), then users will be able to post content with basic HTML tags and the *full HTML* option

Important Links List items Add item Edit menu

**Menu settings**

**Path:**

The path this menu item links to. This can be an internal Drupal path such as node/add or an external URL such as <http://drupal.org>. Enter <front> to link to the front page.

**Menu link title:**

The link text corresponding to this item that should appear in the menu.

**Description:**

The description displayed when hovering over a menu item.

Enabled

Menu items that are not enabled will not be listed in any menu.

Expanded

If selected and this menu item has children, the menu will always appear expanded.

**Parent item:**

The maximum depth for an item and all its children is fixed at 9. Some menu items may not be available as parents if selecting them would exceed this limit.

**Weight:**

Optional. In the menu, the heavier items will sink and the lighter items will be positioned nearer the top.

**Save**

Figure 6: Add new menu item page

'Services', etc, to your site.

Drupal has thousands of contributed modules listed on its site. Each offers a particular plug-and-play feature. A complete list of modules can be obtained from the *Modules* page at [drupal.org/project/Modules](http://drupal.org/project/Modules). A few of the common features are private messages (enable private messaging among your website's users), log-in destination (assign different destination pages to your users based on the user role), buddy list (allows your website users to maintain their own contact/buddy list), etc.

To install a new module on your website, simply download the required module, unzip it, and put it in the modules directory of your Drupal website. Then go to *Administer*→*Site Configuration*→*Modules* page, locate your module, enable it, and the new feature(s) available through this module are good to use. Isn't it simple, easy

will allow all HTML tags in the content. There is another option called PHP Code that can be used to write PHP code—good for advanced users.

Additional block configuration has the following options:

*User-specific visibility settings:* You can specify whether users can control the display of blocks. If allowed, users can enable or disable the blocks when they log into your website from their account section.

*Role specific settings:* Blocks can be enabled for particular user roles on your website.

*Page specific settings:* Here you can restrict the display of blocks to particular pages or you can specify not to display them on some pages. Users who are also good developers can write a PHP code that returns *true* or *false*, to control the block visibility.

Having saved the block, you will find it listed on the Block list page, under the disabled section. Now drag and drop it to the appropriate section and save the configuration, as explained above.

You can choose to change the settings of any block at a later stage by clicking on the 'Configure' link next to it from the Administration Blocks page.

## How to configure and manage menus

Menus are a collection of links used to navigate a website. Drupal has flexible hierarchical menu management systems that allow you to create custom menus in addition to those available by default, or from contributed modules. Every menu automatically creates a block with the same name, which can be placed within any available region on a page. By default, Drupal has out-of-box support for primary and secondary menus.

Let's go to *Administer*→*Site Configuration*→*Menus*. To add a new menu, click on the *Add Menu* tab (Figure 5). Fill in the form presented to give the name and description to the menu and click *Save*. Next, you will be presented a page with the following tabs: *List Items*, *Add Item*, *Edit Menu*. This is where you can now add individual menu items. Click on *Add Item*, fill in the details: *Path* (the URL), *Menu Link Title* (this will appear as a link text), and the description. The parent item will be the menu you created in the prior step (Figure 6). On clicking *submit* you will be presented with the list of items in the current menu (Figure 7). Add a few more links in a similar manner.

To change/remove any menu item, click *edit/remove link* under the operations column on the list page. You can also enable/disable a menu item. Drupal displays only enabled menu items. The default menu items cannot be deleted; they can be enabled or disabled.

Having added all the menu items, you now need to visit the *Block Listing* page described above. Every new menu you create adds a new block on the *Blocks Listing* page, where you can choose to show these blocks in any of the available regions as shown in Figure 3. Just drag your menu block to an appropriate section and save the settings.

**Tip:** Drupal presents a short description with every form field to help you while filling up the data in the form. It has the 'weight' field that allows you to re-order list items, menu items, blocks, apart from many other elements.

## How to give a new look and feel to your website

Your website should reflect your identity! Drupal has a nice architecture. It has the theme layer, also referred to as the 'Presentation Layer' that's responsible for the look and feel of your website without affecting your Drupal core. There are several contributed themes available on the [Drupal.org](http://drupal.org/project/Themes) ([www.drupal.org/project/Themes](http://www.drupal.org/project/Themes)) in *zip* or *tar.gz* format. Download the one that is best suited for your website.

Having downloaded the theme, uncompress it in the *themes* folder of your Drupal install. For example, I downloaded the 'Abarre' theme from [drupal.org/project/abbarre](http://drupal.org/project/abbarre). Next, go to *Administer*→*Site building*→*Themes* page. You will find your new theme 'Abarre', added to the list. Check the checkbox next to it, and choose the radio button next to it. Save the configuration. See how your website has got a facelift now. The default theme was with a three-column layout, while the new one is a two-column layout. Go to the home page and navigate around your website to see how the change in *theme* has affected your website.

You will find that more configuration options are available with this theme. To learn all about that, just go to the *Administer*→*Site Building*→*Themes* page once again, and click the *Configure* link next to the 'Abarre' theme. This theme comes with multiple colour schemes, as well as a custom one to create your own colour scheme. Try changing different colour options and page elements like the logo, site slogan, mission statement, site name, etc.

In this article we discussed how to add new features to your website, how to manage navigation within it and ways to give it a facelift. In our next article we will cover the Drupal directory structure, user access permissions and designing your own content type. 

## REFERENCES:

- Drupal Modules page: <http://drupal.org/project/Modules>
- Drupal Terminologies: <http://drupal.org/node/937>
- Getting Started: <http://drupal.org/getting-started/5/install-contrib>
- Drupal Themes: <http://drupal.org/project/Themes>

**By: Dheeraj Dagliya** is Director in Gloscon Solutions Pvt Ltd, Ahmedabad, a 100 per cent EOU. Gloscon is an Open Source Product Development company with specialisation in Drupal and Ruby on Rails (RoR) frameworks, and has been at the forefront in Drupal development in India. Recently, Gloscon Solutions organised the first ever Drupal Camp in India.



# Who let the geeks out?

IIT Madras and  *L&T Infotech* present  
 **Shaastra 2008**

in association with  and  Reliance



[www.shaastra.org](http://www.shaastra.org)

*Simplify Perfection*

Principal Sponsor  
 **L&T Infotech**

Associate Title  
Sponsors  
 **QUALCOMM**  **Reliance**  
Industries Limited  
Growth is Life

Synergy Associate  


Technical Media  
Partner  


Co-Sponsors



**Vestas**  
Innovative Energy



GE imagination at work



[www.shell.com/careers](http://www.shell.com/careers)

Print Media  
Partner

 **EDUCATION PLUS**

Conclave Magazine  
Partner

 **THE WEEK**  
Your weekly newsletter



# Quingy

## An Alternate Desktop Manager

If building a Linux system without KDE, GNOME, or even X, yet with a nice-looking login method sounds impossible, read on...

**A**nyone using Linux for a while would surely have come across the term 'display manager'. Yes, it is that piece of software that you encounter when you start X on your PC. Display managers (or DM, for short) are the prompts where you enter your

user name and password. It can be used to perform a variety of tasks, including selecting the session that you want to log in to, the language you want, or simply to shutdown or reboot your PC.

There are a wide variety of display managers around with varying levels of

functionality. The two most popular are, of course, KDM and GDM, which come with KDE and GNOME, respectively. They support a number of features, more than most people will ever use. Some of the more convenient features include auto login, support for changing themes, etc. For remote display management, they also have a feature known as XDMCP.

Both KDM and GDM depend on the X server to run. What if you don't want an X server at all? How do you log in then? The command-like login prompt that you use to log in to a virtual terminal is managed by a program called Getty. The command-line login prompt that you use to log in to a virtual terminal is managed by a program called Getty, which is also responsible for invoking the */bin/login* shell to complete the login process.

## Qingy is not Getty

In a world where you have choices, naturally even something as basic as Getty has alternatives. Qingy is one such DM intended to replace Getty. But why would anyone want that? Well, although it can be used from the command line interface, with the help of DirectFB it can provide a nice and fast GUI without the overhead of X server. As expected from a DM, you can log in and start the session of your choice—be it a text console, GNOME, KDE, or whatever else you prefer.

Still don't think it's worth a try? The following are some of the features, as the project homepage at [qingy.sourceforge.net](http://qingy.sourceforge.net) likes to advertise:

- It remembers the last user who logged in, hence focuses on his/her password instead
- Remembers the last session each user chose
- Alternatively, it remembers the last session on a per-virtual console (tty) basis
- It is themable. You can select your favourite theme, or set it to pick a random one every time. A theme will look the same on all machines, independent of the resolution.
- Both text and X sessions are supported, including the ability to start more than one X session at once. You can even start X inside a console when X sessions are already running.
- Additionally, it has support for PAM, screen savers, auto login, session locking and timeout (with lock or logout as available actions), tty-specific options, customisable key bindings, etc.

## Up and running

If that feature list managed to impress you, I'm sure you'd now want to give it a try by installing it. Before we get started, the first step is to install a program called fbset, a utility to check if proper framebuffer video mode is set. You will also need directfb to get a nice GUI.

Now you are ready to install Qingy! For most distributions you can get Qingy and Qingy theme packages from the official repositories using the default package manager only. Just in case it is not available, you can download it from [qingy.sourceforge.net](http://qingy.sourceforge.net).

Once done downloading, untar the Qingy source file, cd to the untarred directory, and run the following command:

```
./configure --PREFIX=/usr --sysconfdir=/etc --localstatedir=/var  
--datadir=/usr/share
```

And in case you want crypto, you can also append the following to the above command:

```
--enable-crypto=none/openssl/libgcrypt
```

If you have missed anything, you will get an error at this point. Note that to compile from source, you will be required to install development packages for ncurses, openssl, xlib, etc.

Once *configure* runs successful, the next step is to run *make*. Follow this by switching over to the root and running *make install*. Assuming that everything went as it was supposed to, Qingy is now installed.

The next step is to remove your current display manager—GDM, KDM, XDM, etc—from starting at boot time. Consult your distribution's documentation on how to do this. Now, we need to make sure that Qingy starts instead. Replace the following line in the */etc/inittab* file:

```
1:2345:respawn:/sbin/getty ttys
```

with

```
1:2345:respawn:/usr/sbin/qingy ttys
```



In case your distribution uses Upstart instead of the old init system, like in Ubuntu, you will have to edit the */etc/event.d/ttys* file instead.

Follow the above step for each tty you want to enable Qingy on. But, remember to leave one tty running Getty to avoid any trouble.

Now either reboot your system or run the following commands:

```
init q  
killall getty
```

That's it; you'll now be greeted by Qingy every time you boot your computer.

You might even want to install extra themes for Qingy at this point. If so, download theme packs from [qingy.sourceforge.net/themes.php](http://qingy.sourceforge.net/themes.php). Installation is as simple as what follows: untar the theme pack; cd to the directory containing the theme pack; copy all the directories from the theme pack to */usr/local/share/qingy/themes*

## Tricks and tweaks

Although Qingy is now ready to use, it's always fun to spice it up a bit by configuring it to our tastes.



Figure 1: A Matrix themed Qingy window manager

First, let's edit the welcome messages that we receive. To do that, open `/etc/qingy/messages` and place your custom messages for each user.

Now we will come to the most important configuration file for Qingy, `/etc/qingy/settings`. In this file you can set all the important options. Some options that you might want to take a look at are:

- `screensaver_timeout`: how many minutes to wait before starting a screensaver
- `screen_powersaving_timeout`: how many minutes to wait before entering power saving mode
- `theme`: the theme to use.
- `last_user_policy`: remember the last user globally, or tty wise
- `last_session_policy`: remember the last login session user wise, or tty wise
- `lock_sessions`: whether to lock the session when tty is switched on or not

You can also have per-`tty` configuration separately in the `/etc/qingy/settings` file by placing the options under the `tty` variable, for example:

```
tty = 2
{
    theme= "matrix"
    screensaver=random
    autologin
    {
        username="mehul"
```



Figure 2: Another Qingy theme

```
    password="mypassword"
    session="icewm"
}

}
tty = 3
{
    theme = random
    screensaver = random
}

}
```

So, now you will have a matrix theme (Figure 1) on `tty2` and will be able to auto-login into the IceWM session. Whereas on `tty3`, you will just get a random theme every time and a login prompt. In case you don't like the framebuffer interface and prefer the Getty-style login prompt instead, you can disable themes but still have all the other features available! 

## REFERENCES:

- [http://en.wikipedia.org/wiki/Getty\\_\(Unix\)](http://en.wikipedia.org/wiki/Getty_(Unix))
- [http://en.wikipedia.org/wiki/X\\_display\\_manager](http://en.wikipedia.org/wiki/X_display_manager)
- <http://qingy.sourceforge.net>

**By: Mehul Ved** is a FOSS enthusiast interested in technology. He is fond of command line and networking.







FREDERICK NORONHA

## Losing That Battle?

The impact of FOSS can be felt everywhere. But are we claiming credit where it's due?

**T**ake three updates that I came across recently: a geek's analysis of Google Chrome, the use of GreaseMonkey for IIT-Delhi, and news from the Khadi Board to go 'paperless' in Kerala.

Check out what Satya Bhat has to say about his experience on running the (much-hyped) Google Chrome browser on GNU/Linux. He tells us he was using Wine for that. What he encountered is explained online at <http://tinyurl.com/6d8ncx>

Bhat appreciates the pluses of Chrome, but adds, "Google Chrome is basically the best of Firefox, Opera, Konqueror, Safari and Internet Explorer, all rolled into one." Which is something that struck me too. There are features that have a haven't-I-seen-it-somewhere-earlier feel to them.

As many of us would agree, Firefox is doing pretty well. It has attained global records over its number of downloads on a particular day.

But still, out of sheer force of habit—aided with some amount of powerful vendor convincing—everyone thinks of the "Big e" as the default browser. The average Net surfer would still use a product like Internet Explorer, despite its many known flaws. Like Colgate once was synonymous with toothpaste (at least in India) and Xerox with photocopying, the Internet Explorer remains the synonym for a browsing tool.

Things don't stop there.

Go to Ishan Arora's page at <http://ishanarora.googlepages.com> and locate his GreaseMonkey script. It's interesting to see how he's using it.

Arora points out that the student registration website at IIT Delhi is really old. As he puts it, it is so old that it has VB code in it. And the site doesn't work without VB support (read: Internet Explorer). So, his GreaseMonkey user script is meant to work around this problem.

At one level, it's interesting to see innovative approaches and workarounds of the *desi* kind. On the other hand, can't those responsible avoid such pitfalls in the first place?

Meanwhile, Anivar Aravind, that determined campaigner from down south, also has some interesting news. He says the Khadi Board hopes to become the first public sector organisation in Kerala to go 'completely paperless' by December 2008. Of course, GNU/Linux will power that.

Swathanthra Malayalam Computing [<http://smc.org.in>]

and its members have been responsible for this, we're told. The Khadi Board is now being considered as one of the FOSS success stories.

We also know of other issues deeply relevant to India. The OLPC (One Laptop Per Child Project) is showing a determination to set up base in our part of the world, despite efforts to thwart it by officials who hardly allowed any discussion on the matter. Perhaps we need to reopen the discussion on how a single statement by a single official could so effectively block the OLPC wheels.

Incidentally, London's *Sunday Times* of August 10, 2008 carried an article titled "Why Microsoft and Intel tried to kill the XO \$100 laptop" [<http://tinyurl.com/5koe2w>].

Some issues brought up here are indeed insightful. But, at the end of the day, one is left to ask: "Do technology options get decided on merit alone? Or is it the lobbying powers of Big Business that decide what options we really have?"

The point is simply this: GNU/Linux is making some rather impressive gains. Its influence is being felt in varied fields. India is contributing its bit too.

If only we could cash in on the resultant publicity that should arise. Unfortunately,

many of these initiatives aren't tom-tommed about much. Leave aside the mainstream media, there's little awareness even in FLOSS circles as to what our victories are.

This leads me to say that while the FLOSS front is doing a good job on the tech side, we're failing in our task of claiming credit.

If we had the cunning and boastfulness of proprietary software, would not the benefit of our work become more appreciated? Or maybe we're just better off without that. Hype is not really needed; but what about simply claiming credit where credit is due? 

**Frederick Noronha** is a Goa-based freelance journalist, who writes on Free/Libre and Open Source Software. The author has co-founded a voluntary, not-for-profit and unfunded group Bytesforall.org. He can be reached at fred@bytesforall.org

GUEST COLUMN

# Programming in Python for Friends and Relations



## Python

# The Glue for IT

If you have a number of programs that need to be run in a sequence and you need an easier programming environment than a shell to control the flow, consider Python. It is widely used by distributions like Fedora and Ubuntu for utilities programming and for graphical front-ends for Linux command line scripts.

**P**ython has an excellent set of built-in modules for interfacing with the OS and the system. So, start the python interpreter and get a list of the modules available:

```
>>>help('modules')
```

Spend a few minutes looking at the long list of modules you can use. You will notice that there are two modules with names that remind us of the work we need to do—'os' and 'sys'. You will want to know what each of these modules does. So, try:

```
>>>help('os')
```

The 'os' module provides OS-dependent functions for managing files, directories and execution of programs. The first method you can try is 'system'.

```
>>> import os
>>> error = os.system('OpenOffice')
sh: openoffice: command not found
>>> print error
32512
```

You will realise that the command for

starting OpenOffice.org is *ooffice* (works on Fedora and Ubuntu) or *soffice*, depending upon the distribution/installation. So, try again but make sure that you close *ooffice* if you are using it:

```
>>> error = os.system('ooffice')
>>> print error
0
```

OpenOffice.org starts, but the control does not go back to the Python program. You will notice that the print statement is executed only after you have closed OpenOffice.org. However, you need to keep the office suite running in the background. An easy way out is to use the '&' option to run the program in the background. So, try the following:

```
>>> error = os.system('ooffice &')
>>> print error
0
```

This time OpenOffice.org starts up and the print statement is executed. You can work in both applications. There remains a slight hitch. You may wish to stop the program. You need to spawn another program and have control over it. You can have a look at the 'spawnlpc' method. It takes a 'no wait' option and returns the process ID. So, try the following commands:

```
>>> pid=os.spawnlpc(os.P_NOWAIT, 'ooffice')
>>> print pid
12058
```

Just to make sure that you have the process ID, try:

```
>>> os.kill(12058,9)
```

OpenOffice.org should close.

Now you can explore the 'sys' module. Two items will be useful from this module. The exit method and the list, argv, which contains the command line parameters passed to the Python program. Write the following code in *test\_params.py*:

```
import sys
import os
for param in sys.argv:
    print 'Parameter = ', param
```

Now, from the command line, run the code below:

```
$ python test_params.py loans_v2.py
Parameter = test_params.py
Parameter = loans_v2.py
```

Notice that the first parameter is the name of the Python script and the other parameters follow. From the

command line, try:

```
$ gedit test_params.py
```

You will find that the editor starts and opens the *test\_params.py* file. So, you can try the same thing from Python:

```
>>> p= os.spawnlpc(os.P_NOWAIT,'gedit', 'test_params.py')
>>> print p
12953
```

The editor starts but does not open the file. As we noticed earlier, Linux expects the first parameter to be the identifier of the application. So, you can pass one additional parameter as follows:

```
>>> p= os.spawnlpc(os.P_NOWAIT,'gedit','gedit', 'test_params.py')
>>> print p
12997
```

The editor has opened the file as expected.

## The workflow

You now have the basic tools at your disposal to start integrating the various tasks from last month's programs. You can start creating a file, *workflow.py*:

```
import os
pid=os.spawnlpc(os.P_NOWAIT,'ooffice','ooffice',
'-accept=socket,host=localhost,port=2002;urp;','-
invisible')
os.system('python db2oo.py')
```

Now, try running this program from the command line:

```
$ python workflow.py
File "db2oo.py", line 18, in <module>
    calc = oo_calc()
File ".... open_office.py", line 14, in oo_calc
    'uno:socket,host=localhost,port=2002;urp;StarOffice.
ComponentContext')
open_office.NoConnectException: Connector : couldn't connect to
socket (Success)
```

You get an error. A little experimentation will lead you to the conclusion that the OpenOffice.org program has started, but is not yet ready for accepting connections. So, you will need to wait before starting the rest of the script. You should revise your *workflow.py* as follows:

```
import os
import time
def start_oo_daemon():
```

```

# Start Open Office in background mode
pid=os.spawnlpe(os.P_NOWAIT,'ooffice','ooffice',
                 '-accept=socket,host=localhost,port=2002;urp;',
                 '-invisible')

print 'Open Office started. Wait a few seconds to connect'
time.sleep(4)

start_oo_daemon()
print "Load Data from db"
os.system('python db2oo.py')
# stop background copies of soffice & soffice.bin
os.system('pkill soffice')
# Manually edit the spreadsheet
os.system('ooffice /home/anil/workbook.ods')
start_oo_daemon()
print "Store data back in db"
os.system('python oo2db.py')
# Clean up
os.system('pkill soffice')

```

You need to run OpenOffice.org in the background twice. So, convert it into a function. A second point to take note of is that the *ooffice* command starts two programs—*soffice* and *soffice.bin*. The PID under consideration is of the first. Unfortunately, the second one is not killed if you have run OpenOffice.org in the background mode and you kill the first one. A simple solution is to use the *pkill* command, as above. By the way, you can leave out the *-invisible* option while developing and testing.

## Using command-line parameters

A well-written program should not depend on hard-coded file names. So, you should accept command line variables. The files needed are the friends and items databases and the temporary spreadsheets workbook. Typically, Linux utilities use the pattern ‘-o value’. You may decide to use ‘-i’ for items, ‘-f’ for friends and ‘-w’ for workbook. So, create a file *test\_params.py*:

```

import sys
print 'The number of arguments',len(sys.argv)
print 'The script name',sys.argv[0]
index=1
while index < len(sys.argv):
    if sys.argv[index] == '-i':
        items_db = sys.argv[index + 1]
        print 'Items ', items_db
    elif sys.argv[index] == '-f':
        friends_db = sys.argv[index + 1]
        print 'Friends ', friends_db
    elif sys.argv[index] == '-w':
        workbook = sys.argv[index + 1]
        print 'Workbook ', workbook
    else:
        print "Syntax Error ", sys.argv[index]
        sys.exit()

```

```

index += 2

```

You should now execute this script, e.g.,

```

$python test_params.py
$python test_params.py -i items.db
$python test_params.py -i items.db -w workbook.ods
$python test_params.py -f friends.db -w workbook.ods -i items.db

```

This script should now be integrated with *workflow.py*. Since it will require changes in the code from last month’s article and is not difficult, you can do this as an exercise.

## Using Tkinter for accepting options

A second exercise worth doing is to not restrict yourself to just command-line parameters. If the parameters have not been given, you could bring up a tkinter form as discussed in an earlier article. You can type the following in *tk\_params.py* and expand it:

```

import Tkinter
import tkFileDialog

def tk_get_params():
    root = Tkinter.Tk()
    items_db = tkFileDialog.askopenfilename(
        parent=root,
        initialdir='/home/anil',
        title='Please select Items DB')
    workbook = tkFileDialog.asksaveasfilename(
        parent=root,
        initialdir='/home/anil',
        title='Please select Worksheet Name')
    return items_db, workbook

print tk_get_params()

```

In the case of *items\_db*, the file must exist; so, ‘askopenfilename’ is the useful method. In the case of the workbook, if the file exists, a warning that an existing file will be overwritten is useful; hence, ‘asksaveasfilename’ is the appropriate method.

As you have tried this, it is easy to add existing widgets and quickly capture the desired parameters conveniently. The hardest part of using *tk* widgets is finding the documentation on how to use them!

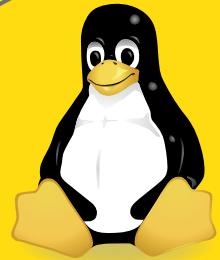
The best place to learn more about various Python modules is by using Python Docs—<http://docs.python.org>. Usage of widgets is best explored by using Google.

You may well think that the world is moving to Web 2.0 so what’s the point of continuing with pre-Web programming. So, next time you will explore how to Web-enable your application to keep track of your loaned items.  

**By: Anil Seth, consultant, [seth.anil@gmail.com](mailto:seth.anil@gmail.com)**



# Makefiles for Your Programs



Here are a few basics, as well as some dos and don'ts while writing makefiles.



The *make* and *build* system is said to be a very important component of any given software project (whether open or closed source). The reason? A mistake in the *make* always results in high costs in development time. Writing simple and sweet, but smart and readable makefiles is an art and can only be achieved with practice.

You can write makefiles manually using *gnu-make* if the project is small, while for bigger projects you can automate things using the *autoconf* and *automake* utilities. However, before we get into that, let's try to understand the concept of *make*.

## Why *make*?

Consider a project development environment where there are hundreds of source files contributed to a given executable. In such projects the compilation times are in hours. Here the developer does not have the luxury of recompiling all the files on-the-fly to make a tiny little change. Here *make* is the real hero.

What *make* maintains is the timestamp information. It maintains the information on whether the given file is updated after the last build or not. Accordingly, when we run *make* again, it will only compile those selective source files that are directly or indirectly modified after the last *make*. This way it avoids the recompilation of thousands of files when only a couple of them have been modified.

## Invisible makefiles

When compiling a single source file, there is no need to write a makefile. For a single source file like, for instance, *MyProg.c*, simply invoke the following:

```
make MyProg
cc MyProg.c -o MyProg
```

So *make* is a command that invokes a makefile, if found. What we have seen just now was the default behaviour of *make* if it does not find a makefile.

## Setting up your text editor

In order to build more than one source file, one needs to have a makefile in place. Therefore, go ahead and create one text file and name it *makefile*. An important thing to remember while editing makefiles is to turn the *Automatic tab expansion feature off* if it's enabled in the text editor. You can do it in *vim* by using '*Esc + :set noexpandtab*'. Existing tabs in a file could be viewed in the vim editor by using the commands '*Esc + :set list*' and '*Esc + :set nolist*'.

## Targets and dependencies

Before we start writing a makefile, let us understand a basic rule that governs *make*. The rule can be classically stated as follows:

```
<target>: [ <dependency> ]*
[ <TAB> <command> <endl> ]+
```

The *target* here is generally what you want to make by using *dependencies* stated. (Note the colon ':' in between.) In the second line we should use a tab (multiple spaces here will not work) followed by the *command* used to make the *target* from the *dependency*.

For example:

```
MyProg.o : MyProg.c header.h
cc -c MyProg.c
```

...or even:

```
Proj.tgz : Module1.c Module2.c Header.h
tar -cvzf Proj.tgz Module1.c Module2.c Header.h
```

Here *MyProg.o* or *Proj.tgz* are the targets; *MyProg.c*, *header.h*, *Module1.c*, *Module2.c* and *Header.h* are dependencies. Note that *make* will always check whether they are present and ensure they are present before attempting to make the targets by executing the commands.

Going further, targets need not be files that really exist. For example:

```
clean:
  rm -fr /home/projpath/*.o
```

Here, the target 'clean' doesn't really exist. It has no dependencies, too. But the action here, which is the removal of .o files at a given path, is important and will be performed when we invoke the target *clean*. (Such targets are classically referred to as phony targets.)

## Local variables

One can define some local variables or so called macros in makefiles to store important information such as utility names, paths, etc, like:

```
CC = gcc
CFLAGS = -Wall -Os -o
PROJPATH = /home/Nilesh/makedemo/
```

...and even:

```
HEADERS = Header1.h Header2.h
MODULE1 = Module1.o
MODULE2 = Module2.o
```

Now we can write:

```
$(MODULE1): $(HEADERS)
  $(CC) $(CFLAGS) $(PROJPATH) /Module1.c
```

...and so on.

## Invoking the targets

As stated before, a makefile has a list of targets and commands to make the targets based on the commands and dependencies. Now users can choose which target to make either from the command line or through the dependencies itself.

For example:

```
make clean
```

...or:

```
Proj: module1 <dependencies for complete project>
      <Set of commands>
```

```
Module1: <dependencies for module1>
      <Set of commands>
```

Now, the following command will let *make* go and check the rule for target 'Proj':

```
make proj
```

Since 'Module1' is listed as a dependency in 'Proj', it will check whether it exists and is 'up-to-date'. This means that if the

dependencies for Module1 are not changed since the last run of *make*, it will first make Module1. In this way, *make* will, one by one, obtain all the dependencies for the target and then make the target 'Proj' itself, in the end. If any of the dependencies for 'Proj' are not found or could not be created, *make* will fail in between and 'Proj' will not be created.

Note that you can write comments in makefiles simply by appending a # (hash mark) before the comment—the same way as it's done in shell scripts.

## Writing a makefile

Now, we are all set to write our first makefile. Let us say we have *file1.c* and *file2.cc* located at */hone/Nilesh/makedemo*. Both of them include *header.h* located at the same path. What we want is an executable called 'prog' out of the two. A makefile will look something like the following code:

```
##### Our First makefile #####
# compiler and related flags
CC = gcc
CFLAGS = -Wall -Os
# Program path
PROGPATH = /hone/Nilesh/makedemo
# Main target
prog : file1.o file2.o
$(CC) $(CFLAGS) -o prog file1.o file2.o
  Echo "Make-prog complete"
#Child targets

file1.o : file1.c header.h
$(CC) $(CFLAGS) -c file1.cc
echo "Compiled file1"

file2.o : file2.c header.h
$(CC) $(CFLAGS) -c file2.cc
echo "Compiled file2"

#To clean the existing .o files.
clean :
  echo "Deleting the stale object files and executables"
  rm -fr $(PROGPATH)/*.o
  rm -fr $(PROGPATH)/prog
```

## Special makefile macros

In the above scenario, we only had two source files. What if I have hundreds of source files? Am I supposed to write a rule for each file? No! The following is a shortcut:

```
%o:%c
$(CC) $(CFLAGS) -c $<
```

Too cryptic? Here '%o' is the target and hence the rule caters to all .o files in the current makefile. '%c' is the dependency and it means the corresponding .c files in the working directory. For example, for *file.o*, the dependency would be *file.c*, etc. The symbol '\$<' means the name of the first dependency, which again means the

source filename itself. Hence, for *file.o*, the rule could be decoded as:

```
file.o:file.c
$(CC) $(CFLAGS) -c file.c
```

The following is the complete list of special macros:

- **\$@** —The file name of the target.
- **\$<** — The name of the first dependency.
- **\$\*** — The part of a filename that matched a suffix rule.
- **\$?** — The names of all the dependencies newer than the target separated by spaces.
- **\$^** — The names of all the dependencies separated by spaces, but with duplicate names removed.
- **\$+** —The names of all the dependencies separated by spaces with duplicate names included and in the same order as in the rule.

Refer to [www.cprogramming.com/tutorial/makefiles\\_continued.html](http://www.cprogramming.com/tutorial/makefiles_continued.html) for documentation on how to use them.

## Auto generating the dependencies

Consider a file such as *storage.h*:

```
*****storage.c*****
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>
#include "sas_drv.c"
#include "sata_drv.c"
:
:
```

Now, let's assume *sas\_drv.h* looks something like the following code:

```
*****sas_drv.h*****
#include <stddefs.h>
#include "basetypes.h"
#include "sas_ssp.h"
#include "sas_smp.h"
:
:
```

We have nested dependencies here. A source file depends on a header file, which in turn depends on multiple header files. Writing the dependency rules here manually is really difficult. Here, a tool *makedepend* comes in handy.

The rule looks like what's shown below:

```
DEPEND = makedepend
SRCS= file1.c file2.c file3.c

dep:
$(DEP) $(SRCS)
```

and *makedepend* will scan through all the source files and append the current makefile with the appropriate dependency rules.

## Supermake

Consider a project in development where there are multiple modules. Each module has a makefile and now we are supposed to 'make' one or more modules selectively depending on requirements and put them together in a given binary. This can be achieved by writing a different makefile for the different modules in their corresponding paths and writing a 'supermake' file that will invoke the makefiles by going in to selected subdirectories. The makefile will look like what's shown below:

```
MAKE = make
PATH_MODULE1 =/home/Nilesh/proj/module1
PATH_MODULE2 =/home/Nilesh/proj/module2

target1: module1 module2
        #Commands here to make the 'target' from module1 and module
2.

module1:
$(MAKE) -c $(PATH_MODULE1)
#And make will invoke the makefile at PATH_MODULE1 if found

module2:
$(MAKE) -c $(PATH_MODULE2)
#And make will invoke the makefile at PATH_MODULE2 if found
```

## Not only make

Consider the following makefile rules:

```
release:
$(TAR) -cvzf Proj.tgz $(SRCDIR)/*.* $(INCLDIR)/*.*.h
$(MV) Proj.tgz $(RELEASEDIR)
Echo "Release made.. please find the tarball at Release location"
```

or:

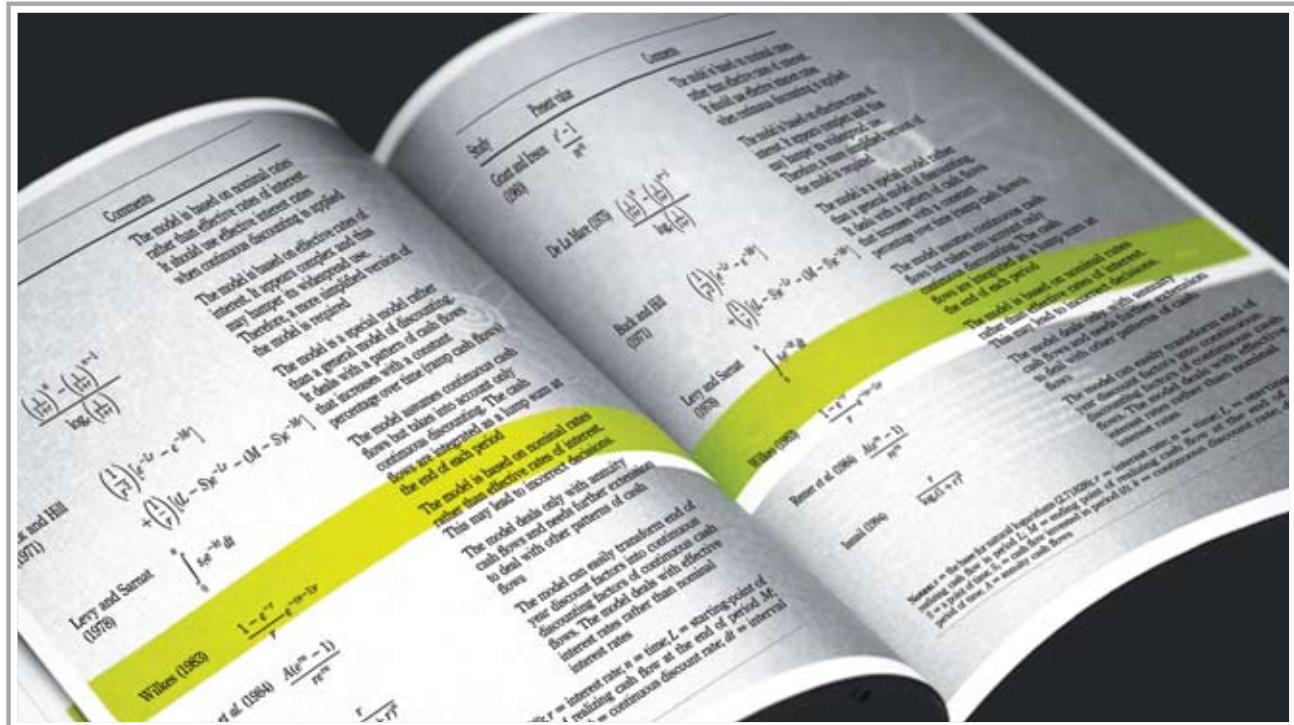
```
archive:
$(AR) -rc $(LIBRARY) $(OBJS)
$(MV) $(LIBRARY) $(LIBDIR) /
Echo "Library built and moved at load location"
```

Apart from merely using the command line, you can do many other things with the *make* framework, and do so with more elegance.

## Further reading

This article throws light on the basic syntax and semantics of the makefiles and some of the advanced features. Your projects can be better managed if the makefiles are written correctly and concisely. There is a lot more that could be found at [www.gnu.org/software/make/manual/make.html](http://www.gnu.org/software/make/manual/make.html). The question is, are you ready to try the 'make' recipe? 

**By: Nilesh Govande.** The author is a Linux enthusiast and could be contacted at [nileshgovande@yahoo.com](mailto:nileshgovande@yahoo.com). His areas of interest include Linux system software, application development and virtualisation. He is currently working with LSI Research & Development Centre, Pune.



# Creating Beautiful Documents

# Using LaTeX

If you care about creating polished and stylish documents, and if you have a sense of aesthetics and like some finesse, try LaTeX! Believe us, you'll most likely fall in love with it and use it for the rest of your life!



aTeX is a document preparation system. It is used for high-quality typesetting and with it, you can create beautiful, professional-looking documents. LaTeX is free and is distributed under the LPPL (Latex Project Public License), which is slightly different from the GPL.

Who uses LaTeX? Well, LaTeX is used by small publishing companies and large software organisations; students and Nobel prize-winning researchers; designers creating artistic posters, as well as authors writing journal papers. In other words, it is used worldwide by people who care about writing ‘quality’ documents.

There are major differences between the word processing we are used to and what's available in LaTeX. Unlike the formatting we do with the content of the document in WYSIWYG (what you see is what you get) word processors, creating documents in LaTeX is more like writing computer programs—we write LaTeX code and 'compile' it (yes, we compile the code!) and generate the document in the desired format. More technically, LaTeX is more like a 'mark-up language', just like XML or HTML. Unlike WYSIWYG word processors where content is mixed with formatting, in LaTeX, content of the document is separated from formatting or the display mechanism (a rough analogy is using `.css` files for `.html`):

so, we can concentrate on the content of the document and leave the specific formatting details to the software.

At first, I was totally surprised that we needed to write code (like we do in programming) to create documents, and LaTeX can be slightly intimidating in the beginning. It is similar to the experience I had with the vi editor: I was used to programming with interactive editing and at first I was surprised that there were two different modes that I had to work with. Once I got used to vi (or vim) for programming, I didn't like anything else, though it took time to learn the commands and get used to it. Similarly, I found it surprising that I should write code for LaTeX, but then after using it for some time, I've got used to it and now I only enjoy writing in it.

## Why LaTeX?

The answer is simple: if you want to write documents that just look great then the best option is to use LaTeX. Sometime, you'll be forced to learn LaTeX. If, for instance, you're doing your M. Tech or PhD and you want to write a paper or thesis, more likely than not, you'll be forced to create your document in LaTeX. For some niche domains like mathematics, physics or chemistry where lots of symbols, complex equations, tables and cross-references are involved, there is rarely any choice but to use LaTeX. Although mostly used in academic circles and in commercial publishing, don't underestimate the popularity and wide use of LaTeX—even novels and short stories are occasionally written using LaTeX!

Those who don't care much about writing good-looking documents—I call them 'impatient' writers—and those not from computer science or academic background, prefer using interactive word processors. Novices/beginners who don't want to write LaTeX code can use interactive document processors that create LaTeX output such as LyX. But remember: if you want to make the best out of LaTeX, writing code is the best way to go.

It's best to show the difference visually. Just check out Figures 1 and 2: the document is written in MS Word in the former, while the other is generated with LaTeX. Those who don't know anything about typesetting can also find minute differences between two documents: the spacing between the words, how italics look, how the superscript and subscript symbols look, how the default settings look (these two documents were written with default settings on—and no, in MS Word, the bold face was not enabled for the text, but it looks dark; so to differentiate the title, I've used bold there). To avoid bias, I've used the same fonts—for MS Word, I used the default Times New Roman font; for LaTeX I've used the *pstlatex* package that replaces the default computer modern fonts with PostScript fonts.

## LaTeX advantages

The first and foremost is the *quality of output*. In my experience, LaTeX produces the best quality output and I would recommend you experience it yourself.

Avoid frustrations in creating large documents. Anyone

### Short History of L<sup>A</sup>T<sub>E</sub>X

September 8, 2008

For those who keep wondering what *TEX* or *L<sup>A</sup>T<sub>E</sub>X* is and why it was developed, here is a short history.

Donald E Knuth, when he was working on his monumental book *The Art of Programming* was not happy with the type-setting for his book. So he created a type-setting system that would let him concentrate on the content of the book rather than getting the type-setting software distracting and getting him worried about formatting the output. For example, instead of doing formatting for a quotation every time with center aligned and extra border from the page, with different font, italicized, surrounded with double quote etc., he wanted just to say - this is a quote - and wanted the typesetting system to take care of formatting it accordingly. He designed *T<sub>E</sub>X* to do that. However, *T<sub>E</sub>X* was still low level and difficult for beginners to use. So, Leslie Lamport created a set of high level *T<sub>E</sub>X*-macros and made the common tasks - such as creating table of contents - easier. Now-a-days, most of the people use *L<sup>A</sup>T<sub>E</sub>X* instead *T<sub>E</sub>X*.

Figure 1: Default text rendering in MS Word

### Short History of L<sup>A</sup>T<sub>E</sub>X

September 8, 2008

For those who keep wondering what *TEX* or *L<sup>A</sup>T<sub>E</sub>X* is and why it was developed, here is a short history.

Donald E Knuth, when he was working on his monumental book *The Art of Programming* was not happy with the type-setting for his book. So he created a type-setting system that would let him concentrate on the content of the book rather than getting the type-setting software distracting and getting him worried about formatting the output. For example, instead of doing formatting for a quotation every time with center aligned and extra border from the page, with different font, italicized, surrounded with double quote etc., he wanted just to say - this is a quote - and wanted the typesetting system to take care of formatting it accordingly. He designed *T<sub>E</sub>X* to do that. However, *T<sub>E</sub>X* was still low level and difficult for beginners to use. So, Leslie Lamport created a set of high level *T<sub>E</sub>X*-macros and made the common tasks - such as creating table of contents - easier. Now-a-days, most of the people use *L<sup>A</sup>T<sub>E</sub>X* instead *T<sub>E</sub>X*.

Figure 2: Default text rendering in document generated using LaTeX

with experience in using software like MS Word for large documents such as a book will understand how frustrating it is to work with a word-processor for a long time to get formatting, indexing, cross-referencing, etc, correct and consistent. LaTeX is very convenient to use in medium- to large-size documents that are complex. In my experience, for short documents such as memos or letters that we want to get done with quickly, it is not much use.

The next point is *availability*. LaTeX code is portable: you can create a LaTeX document and use in any platform. It is available from AmigaOS, Mac OSX, Linux and Windows. This advantage is not available in many proprietary formats such as the *.doc* format. Lots of free and open source software is available. Since it is used widely, supporting tools, documentation and help is available.

Next in line is the *choice of output formats*. LaTeX is also useful for creating a wide range of document output formats. It is enough to write a LaTeX document (a text file) and, with that, we can create the document in almost any format that we usually use, including *.ps*, *.doc*, *.html* and *.pdf* formats. So we can just keep the text file and generate the desired output format in the given platform,

## Hello world!

```
$ dvips hello
$ gv hello.ps
$ ps2pdf hello.ps
```

And now you can use a PDF viewer (such as Evince or KPDF) and view the PDF file.

If you feel it's a lot of work, you can use any of the LaTeX editors that let you do this work automatically for you.

Figure 3: LaTeX output file

as and when required.

This is very helpful when the content is of moderate or large size and we need different output formats. For example, a publishing house might get articles from its contributors in *.doc* format; typically, they convert it into LaTeX format! With this, they can enforce the common style they use for the magazine for every article, with ease. They can create *.html* files automatically for posting select articles in their websites. They can provide the digital version of the article as *.pdf* versions online for select subscribers. They can create *.ps* (PostScript) files that are print-ready before printing the content. All this can be done with LaTeX documents when we have the necessary software supporting these features.

### LaTeX 'Hello world' code

Let's explore coding with LaTeX with *hello world* code.

```
\documentclass[12pt]{article}
\begin{document}
\centering{ \textbf{Hello world!}}
\end{document}
```

In LaTeX, all the commands start with '\'. The required arguments for the commands are provided within curly braces ({{}}) and optional arguments are in square braces ([]).

The `\documentclass` command is to tell the type of the document, for example, an article, book or letter; here we have selected the article type. We can also mention the size of the font for the document; 12pt is what we've selected here. We can put all our code for the document within `\begin{document}` and `\end{document}`.

In the body of the document, we have introduced `\centering` environment—the content provided within this environment will be page-centred. The `\textbf` command makes the text inside it bold face. That's it; so how do we run it and see the output?

Save it as, *hello.tex*, for instance, and assuming that you already have *latex* pre-installed in your Linux machine, type the following commands:

```
$ latex hello
```

This generates the *.dvi* file that you can view; Figure 3 shows the output. If you want to convert it to a PDF file, type the following three commands:

### LaTeX in Linux

There is a lot of free/open source software available for LaTeX. Many of the Linux distributions have LaTeX software readily available with their installations. We'll cover some of the important and more widely used software here.

**LyX:** It is an interactive document processor that is similar to conventional word processors; so novice LaTeX users will find it comfortable using this. You can download LyX and learn more about it from [www.lyx.org](http://www.lyx.org).

**teTeX:** teTeX is a free TeX distribution system that is widely used in UNIX/Linux platforms. However, it is not actively maintained any more. See [www.tug.org/tetex](http://www.tug.org/tetex)

**Kile:** Kile is an integrated LaTeX environment; it is a friendly text editor that works on KDE. Kile has many useful features: we can use readily available document templates and wizards, compile, convert, and view the documents at one go, use the auto-completion feature to avoid typing LaTeX commands, get more information on commands using context-sensitive help and so on. You can get it from [kile.sourceforge.net/download.php](http://kile.sourceforge.net/download.php)

**Texmaker:** Texmaker is a free and open source LaTeX editor available for Linux (it is also available for other platforms such as Windows and MacOS). See [www.xmath.net/texmaker/index.html](http://xmath.net/texmaker/index.html)



### REFERENCES:

- At the project website for LaTeX, you can download software for LaTeX, read documentation and learn more about it: [www.latex-project.org](http://www.latex-project.org)
- Homepage for the organisation of TeX users known as TUG (TeX User Group): [www.tug.org](http://www.tug.org)
- The Comprehensive TeX Archive Network (CTAN), where we can get LaTeX packages, LaTeX software downloads and much more. When you have specific needs while using LaTeX and need a package to solve a particular problem, this is the first place to look: [tug.ctan.org](http://tug.ctan.org)
- A free book on LaTeX written with beginners in mind; provides a good introduction and overview of the features of LaTeX: [www.maths.tcd.ie/~7Edwilkins/LaTeXPrimer](http://www.maths.tcd.ie/~7Edwilkins/LaTeXPrimer)

**By: S.G. Ganesh** is a research engineer in Siemens (Corporate Technology). His latest book, "60 Tips on Object Oriented Programming", was published by Tata McGraw-Hill in December last year. You can reach him at [sogganesh@gmail.com](mailto:sogganesh@gmail.com)

# LINUX JOBS

**Post: Tech Lead - J2EE**

Company: Ismart Panache Inc

Profile: Should have an experience in building/managing team up to 4-5 people with thorough understanding of Software development life cycle.

Exp: 5-7

Location: Gurgaon

Email: adeshpande@ismartpanache.com

**Post: JAVA technical leads**

Company: Ness Technologies

Profile: 6+ Years Server-side Java development and over 7+ Years IT experience having expert level knowledge and hands-on experience in Java, JDBC, Servlets, JSP, EJB, Weblogic, XML, XSLT, SOAP etc.

Exp: 6-11, Location: Mumbai

Email: karthik.shreeram@in.ness.com

**Post: J2EE Developer**

Company: Oracle India Pvt. Ltd.

Profile: Knowledge and experience of Oracle products preferable with knowledge and experience with Unix, Linux, Java/J2EE, experience in programming using Perl, shell, C, Java or similar languages.

Exp: 8-13

Location: Bengaluru

Email: Nidhi.Jain@oracle.com

**Post: Sr. Software Engineer - Java, J2ee**

Company: CollabNet Software Pvt. Ltd.

Profile: Should have a Bachelors degree in Engg or MCA with 6+ years of development experience in building enterprise grade applications/product, 2+ years of development experience in Unix/Linux/Solaris environment.

Exp: 6-8, Location: Chennai

Email: recruitment@collab.net

**Post: JCAs Technical Position**

Company: Pantaloons Retail India Ltd.

Profile: Needs minimum 4yrs of strong experience in Sun SeeBeyond JCAs or ICAN, Unix (Linux, Solaris, or AIX) experience as development and production environment platform.

Exp: 7-10

Location: Ahmedabad

Email: ritesh.devadiga@pantaloons.com

**Post: Technical Support Manager**

Company: CCS Computers Pvt. Ltd.

Profile: Skills required are as follows: Project Management Skills, ITIL Skills, Good Communication skills, Presentation Skills etc.

Exp: 5-10, Location: Delhi

Email: rakesh@ccscomputers.co.in

**Post: Sr. PV sql Development Professionals**

Company: Agilis Information Technologies Intl. Pvt. Ltd.

Profile: The candidate should be good in PV sql Development. Should have hand on experience on to coding related to oracle (Procedures, functions, triggers etc). Exposure to Unix, linux is mandatory.

Exp: 4-9

Location: Delhi/NCR

Email: careers@agilisinternational.com

**Post: QA Lead**

Company: Agilis Information Technologies Intl. Pvt. Ltd.

Profile: On hands experience in testing of products or modules in Test Lead Role, on hand experience with querying the Oracle database. Long or short term Programming experience.

Exp: 4-7

Location: Delhi/NCR

Email: careers@agilisinternational.com

**Post: ETL Architect**

Company: Careernet Technologies Pvt. Ltd.

Profile: At least 4-5 experience in any RDBMS (preferably TERADATA and good exposure on Oracle), must have hands on (current) experience on UNIX/Linux environment.

Exp: 7-12

Location: Bengaluru

Email: aishwarya.d@careernet.co.in

**Post: Associate Project Manager**

Company: Dexterity Business Analysts Pvt. Ltd.

Profile: The person should have a M.Sc/Ph.D. in Statistics, Econometrics or Mathematics with a solid understanding of statistics with competency in databases (SQL, Oracle), programming languages (C, S, Java), operating systems (Windows (2000, 2003), UNIX, Linux).

Exp: 5-7

Location: Chennai

Email: mohamed.u@dexterity.in

**Post: Sr. Java/J2EE**

Company: BPA Technologies Inc

Profile: This position requires 5-8 years software development experience using standard web technologies, e.g. Java, JSP, SOAP, XML, Hibernate, JBoss, HTML, CSS, Struts. Must have Linux/Unix experience

Exp: 5-8

Location: Visakhapatnam

Email: seshu@bpatech.com

**Post: Performance Tuning**

Company: Sierra Atlantic Software Services Ltd.

Profile: Should have minimum 5 yrs as an Oracle DBA/Architect with demonstrated experience with Oracle 9i and 10g. Experience in working with multiple operating systems (Linux/UNIX/Windows).

Exp: 5-7

Location: Hyderabad

Email: hotjobs@sierraatlantic.com

**Post: Oracle Apps DBA**

Company: Bahwan Cyber Tek

Profile: Should have 3+ yrs of relevant experience, must have good exposure on OS as well as Instance maintenance with an experience in atleast 2 Implementation projects, AIM Methodology.

Exp: 4-9

Location: Chennai

Email: careers@bahwancybertek.com

**Post: .NET 2.0 (Dot Net)**

Company: I Sense India OSL

Profile: Substantial knowledge of UNIX/Linux or Windows environments with an extensive experience in designing and implementing large-scale and small scale systems.

Exp: 5-8

Location: Bengaluru

Email: jobs@indiaosl.nl, s.rego@indiaosl.nl

**Post: Sr. Developer**

Company: Ismart Panache Inc

Profile: The person would be involved in development and maintenance of database (Sybase) with Solaris, hands on development experience in C/C++, experience of working in Sybase, Oracle experience etc.

Exp: 4-4

Location: Delhi/NCR

Email: ptiwari@ismartpanache.com

**Post: Unix System Administrator**

Company: Cyberthink Infotech Pvt. Ltd.

Profile: Experience on Sun Solaris (2.7, 2.9), Veritas Volume Manager, Veritas netbackup with an experience NFS, NIS, TCP/IP, SAN, IBM SVC, McData Switches &amp; Directors.

Exp: 5-10

Location: United States

Email: jobs@cyberthinkinfotech.com

**Post: Sr. Analyst**

Company: BA Continuum Solutions Pvt. Ltd.

Profile: Mandatory skills: C, C++, database, windows or Linux platform experience, experience in developing and designing complex multi threaded applications is required.

Exp: 5-7

Location: Mumbai

Email: paramveers.narang@bankofamerica.com

**Post: Architect C, C++**

Company: IBS Software Services Pvt.Ltd.

Profile: Technical architect experience, experience in developing C/C++ products on UNIX/Linux platform. Should possess Bachelors or Master's degree in IT/Comp Science or related field.

Exp: 7-12

Location: Kochi

Email: careers@ibspcl.com

**Post: Sr Lead- Network Processor Development**

Company: Global Logic

Profile: Good communication skills and experience in dealing with customers with an experience in developing Networking products &amp; exposure to Network Processor Development.

Exp: 6-11

Location: Noida

Email: jashan.singh@globallogic.com

**Post: Build and Release Engineer**

Company: Tanla Solutions Ltd.

Profile: Must have a minimum of 3years of SCM (Software Control Management), preferably with open source CVS and Subversion Experienced with Unix OS with emphasis in Linux environment.

Exp: 4-7

Location: Hyderabad

Email: sharmila.ankireddy@tanlasolutions.com

**Post: Software engineer- C, C++, linux/unix**

Company: Synopsys

Profile: Must have BS in CS with 5+ years of relevant experience or MS with 3+ years of experience in C, C++ programming software for operating systems with prior knowledge experience with multiple Unix operating systems.

Exp: 5-7

Location: Bengaluru

Email: swaroopa@synopsys.com

**Get the Best Paying Jobs @ naukri.com**  
**2,00,000 Jobs to choose from!**

35,000 Recruiters | Get Jobs by email | Block Your Current Employers

Log on to [www.naukri.com](http://www.naukri.com)

Powered by

 naukri.com<sup>®</sup>

India's No 1 Job Site

# CELEBRATE Diwali

## Win EXCITING PRIZES

### Subscribe Now!



**Cover Price :**  
Rs 60/- with CD  
Rs 40/- without CD\*

A magazine for the electronics fraternity... professionals, businessmen and hobbyists

\* From September 2008 issue



**Cover Price :**  
Rs 100/- with CD/DVDs

The complete magazine on open source



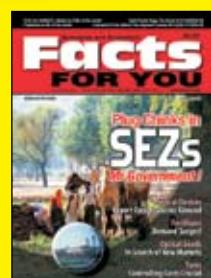
**Cover Price : Rs 50/-**  
An IT magazine for businessmen



**Cover Price : Rs 50/-**  
India's first Electronics Sourcing magazine



**Cover Price : Rs 40/-**  
Information Technology—Today. Tomorrow.



**Cover Price : Rs 100/-**  
India's most influential monthly on Economic and Business Affairs

*Send the filled-in form (Opposite page) or its photocopy to :*



**EFY Enterprises Pvt Ltd, D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020 Ph: 011-26810601-03;**  
**Fax: 011-26817563, E-mail: subscribe@efyindia.com, website: www.efyindia.com**

• Prizes will be selected through Lucky Draw. Gift Items photographs are indicative only. Condition Apply. • These rates are applicable for new subscribers as well as renewal by existing subscribers. • The rates are valid for subscribers within India only. • Please allow 4-6 weeks for processing of your subscription. • Disputes, if any, are subject to exclusive jurisdiction of competent courts and forums in Delhi/New Delhi only.

**1st  
prize**

**LCD TV**

**2nd  
prize**

**MP4 Player  
(10 Nos)**

**3rd  
prize**

**TRAVEL BAG  
(100 Nos)**



**ASSURED  
FREE GIFT**  
with each subscription  
**Round-Neck  
T-Shirt**

Magazine	Unit Price (Rs.)	7 Years (Rs)	5 Years (Rs)	3 Years (Rs)	2 Years (Rs)	1 Years (Rs)
Electronics For You (EFY)	40/-	1680 <input type="checkbox"/>	1445 <input type="checkbox"/>	1015 <input type="checkbox"/>	725 <input type="checkbox"/>	385 <input type="checkbox"/>
EFY (with CD)	60/-	2520 <input type="checkbox"/>	2160 <input type="checkbox"/>	1510 <input type="checkbox"/>	1080 <input type="checkbox"/>	575 <input type="checkbox"/>
Electronics Bazaar	50/-	2100 <input type="checkbox"/>	---- <input type="checkbox"/>	1260 <input type="checkbox"/>	900 <input type="checkbox"/>	480 <input type="checkbox"/>
Linux For You (with CD & DVD)	100/-	4200 <input type="checkbox"/>	3600 <input type="checkbox"/>	2520 <input type="checkbox"/>	1800 <input type="checkbox"/>	960 <input type="checkbox"/>
Information Technology	40/-	1675 <input type="checkbox"/>	1440 <input type="checkbox"/>	1010 <input type="checkbox"/>	720 <input type="checkbox"/>	380 <input type="checkbox"/>
BenefIT	50/-	2095 <input type="checkbox"/>	---- <input type="checkbox"/>	1250 <input type="checkbox"/>	895 <input type="checkbox"/>	475 <input type="checkbox"/>
Facts For You	100/-	4190 <input type="checkbox"/>	3600 <input type="checkbox"/>	2520 <input type="checkbox"/>	1800 <input type="checkbox"/>	960 <input type="checkbox"/>

- Don't Miss A Single Issue!**
- Ensure Regular Supply**
- FREE Replacement Policy\***
- Get Your Money Back if Not Satisfied\*\***
- Order Now For FREE Home Delivery**

\* Replacement will be made upon receipt of the damaged / non-receipt copies within 30 days of its publication

\*\* After three month, if you are not satisfied with the magazine, your balance amount will be returned.

Name \_\_\_\_\_ Company Name \_\_\_\_\_ Designation \_\_\_\_\_

Mailing Address \_\_\_\_\_ Distt. \_\_\_\_\_

City \_\_\_\_\_ Pin Code \_\_\_\_\_ State \_\_\_\_\_ Phone \_\_\_\_\_

Email \_\_\_\_\_ Subscription No. (for existing subscribers only) \_\_\_\_\_

I would like to subscribe to the above ()marked magazine(s), for the period \_\_\_\_\_ Please find enclosed a sum of Rs \_\_\_\_\_ by DD/MO/crossed cheque\* bearing the No. \_\_\_\_\_ dt. \_\_\_\_\_ in favour of **EFY Enterprises Pvt Ltd, payable at Delhi.** (\*Please add Rs 50 on non-metro cheque)

or charge my credit card  VISA  MasterCard Please charge Rs. \_\_\_\_\_

against my credit card No. **LLLLLLLLLLLLLLLL** C V V No. \_\_\_\_\_ (Mandatory)

Date of Birth \_\_\_\_\_ / \_\_\_\_\_ / \_\_\_\_\_ (dd/mm/yy) Card Expiry Date \_\_\_\_\_ / \_\_\_\_\_ (mm/yyyy)

Signature (as on the card)



# In Pursuit of Freedom!

Of late, OSSCube has been quite active in organising unconferences and barcamps to promote and discuss FOSS solutions. There is, of course, a lot more to the company...

**O**SSCube is a niche open source company. It is dedicated to providing OSS solutions to businesses. Started in 2006 by a group of friends—Sonali Minocha, Vineet Agrawal and Lavanya Rastogi—the firm is into open source product development, bespoke development, open source product customisation, consultancy and training. With a team of over 100 members, the company has offices at Noida and Bangalore, with a presence in the US and France too.

The firm has partnered with leading technology development companies, such as Zend, MySQL, Ingres, Moodlerooms and Continuent, to provide many effective solutions. Some of the solutions developed by the firm include enterprise products, BPMN-compliant workflow applications, CCHIT-compliant healthcare applications, Web 2.0 applications, social networking sites, e-commerce applications, as well as widgets, games and mobile applications.

Besides offering services globally, OSSCube is active in India. Here, the firm works with organisations like Yahoo, Reliance, Bigadda.com, naukri.com, Google, Intel and Amazon.

## The aim

“We observed a latent need for quality open source solutions and service. We felt that there was a lack of trained professionals in the open source domain, and the opportunities available to the professionals were limited. This we grasped as an opportunity. The motive behind the inception of the firm was to provide a platform for open source expression and creativity,” reveals Sonali Minocha, co-founder, OSSCube.

The goal of OSSCube is to empower its partners and customers with quality open source solutions by offering an integrated value chain of services. It also aims to increase the extent of open source adoption in India through consistent innovation. The

team finds it surprising that despite having one of the highest skilled technical manpower, India lags behind in its original open source contribution.

## A growing network

"We serve organisations ranging from the entrepreneurial start-ups to the Fortune 500, which span 21 countries. It doesn't matter to us where the customers come from so long as they have the open source itch or have a business case for adopting open source," states Minocha.

OSSCube generates approximately 35 per cent of its revenue from India. "We are bound by a non-disclosure agreement of a potential VC not to disclose these figures. We are a debt-free, profitable company," retorts Minocha when asked about the turnover of the company.

## Spreading the word

Open source is no longer a fad or a pursuit of the geeks. The phenomenal success of companies like MySQL (being acquired by Sun for \$1 billion) and the support of heavy weights like IBM for open source, have put it right in the reckoning amongst strategic choices for the CIOs. Even Microsoft is going partly open source! Minocha opines, "The question is no longer of awareness regarding what open source is all about or if it is out there. The issue that we must focus on is how to create awareness about the products and services that are available in the market. This can only be accomplished by community engagement, as well as end user education. We also need to lobby for support from policy-making institutions like the government of India, just as the other large platform vendors like Microsoft, etc, have done. The open source market is very fragmented, and a large specialised player is yet to emerge in India."

This, in part, is a huge opportunity and responsibility for companies like OSSCube. To engage the community, the firm is focusing a lot on organising unconferences and supporting barcamps across the country, like the OSScamp series of unconferences, designed to heighten awareness about FOSS. The firm also intends to develop strong academic partnerships in the future, enabling it to increase the supply of qualified professionals for the open source engagements. This, in Minocha's opinion, has been the biggest hurdle for both technology as well as end user companies. Besides OSScamp, the company regularly conducts free seminars and road shows in various cities to spread awareness about open source technologies. It also organises free monthly open source meet-ups/ trainings at its campuses in Delhi and Bangalore.

The team at OSSCube believes that it is really important for them to be involved with the open source world at all levels. This starts from the grassroots level and extends well into the global scenario. There is so much to learn in, and from, open source that this interaction is very essential for any organisation in the open source domain. It has taken initiatives to launch

two open source projects at SourceForge, while another one is slated to be launched this year. The firm regularly participates in global open source conferences and shares its best practices, and invites people from the global community to India to participate in OSScamps.

## Constantly evolving

"We are an open source company. We deal only in open source. We earn our bread, butter, jam, and juice from open source. A rise in the demand for open source solutions is good for our business," states Minocha, with a grin. But she is also sceptical because open source is subject to as many threats as the opportunities it creates. This is because with competing open source solutions, each having their own advantages, the firm has to create a skilled workforce. Thus, with increasing solutions, even if the range of the company's business doesn't increase, the service pipeline broadens and loses edge. This requires constant tuning and alignment with the industry and the firm's objectives.

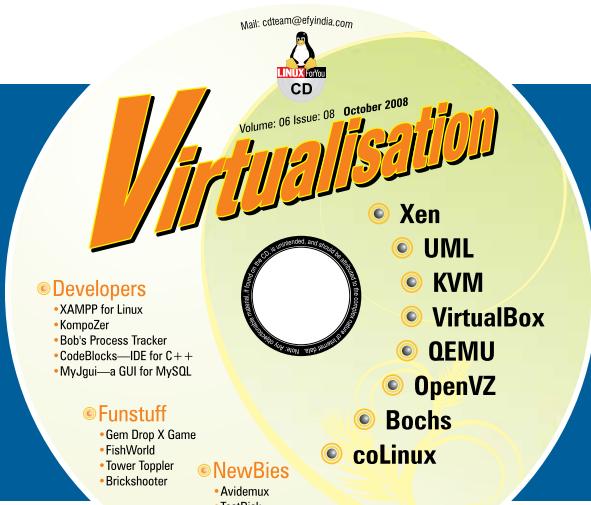
"We foresee the future as a streamlined, interdependent world, where technologies collaborate... In other words, the future is about the interoperability of technologies, wherein you can substitute one technology with another at any given point in time without disrupting the whole mechanism. It's a way of abstraction. The end product would become technology independent," affirms the co-founder.

Minocha states that the team plans to keep doing the right things and improve on what they aren't yet good at. They are engaged in streamlining operations and expanding their capabilities for large open source technology-based product development and remote DBA services. In the pipeline are some open source community-driven projects too, through which the team plans on further engaging the community and creating utilities that would empower any Internet-driven company. But she feels that it is pretty early to be talking about those plans.

## Sound advice

Minocha recommends 'the freedom' to everyone. Her advice to naïve entrepreneurs is to opt for open source platforms. She states, "Just look at what entrepreneurs have accomplished using open source both in terms of ROI and innovation. Before you decide otherwise, at least come up with five good reasons why you should NOT have an open source platform. I think, these days, entrepreneurs are smart and making informed choices. If you look at the Internet space, today it's difficult to find a company that is not using open source. So I think all nascent techno-preneurs need to take a hard look at open source." 

**By: Cholena Deb, LFY Bureau**



# The Virtualisation Bandwagon

Do you have too much machine idle time and wonder if you could put it to some good use? It's time to take the virtualisation drive.

**O**ne year on, with the October 2008 issue of LFY again focussing on virtualisation, we've packed the CD with eight of the best FOSS solutions available.

**Xen 3.3:** The Xen hypervisor acts as a thin layer between the hardware and the operating system, allowing several guest operating systems to be executed on the same computer hardware at the same time.

</software/virtual/xen/>

**VirtualBox 2.0.2:** VirtualBox is a general-purpose full virtualiser for x86 hardware and is targeted for server, desktop and embedded use. It is installed on an existing host operating system. Within this application, additional operating systems (guest OS) can be loaded and run, each with its own virtual environment.

</software/virtual/virtualbox/>

**QEMU 0.9.1:** QEMU is a fast processor emulator. Using dynamic translation it achieves a reasonable speed while being easy to port to new host CPUs. In its user-mode emulation mode, it can launch Linux processes compiled for one CPU architecture, on other architecture. In its full system emulation mode, QEMU emulates a

full system, including a processor and various peripherals.

</software/virtual/qemu/>

**UML 2.6.22:** The User-Mode Linux (UML) is a safe and secure way of running Linux versions and Linux processes. Run buggy software, experiment with new Linux kernels or distributions, and poke around in the internals of Linux, all without risking your main Linux set-up.

</software/virtual/uml/>

**Bochs 2.3.7:** Bochs is a highly portable open source x86 PC emulator that runs on most popular platforms. It is capable of running most operating systems inside the emulation, including Linux, DOS, Windows 95/98 and Windows NT/2000/XP or Windows Vista.

</software/virtual/bochs/>

**KVM (Kernel-based Virtual Machine):** This is a full virtualisation solution for Linux on x86 hardware containing virtualisation extensions (Intel VT or AMD-V). It consists of a loadable kernel module, *kvm.ko* that provides the core virtualisation infrastructure and a processor specific module, *kvm-intel.ko* or *kvm-amd.ko*. Using KVM, one can run multiple virtual machines with unmodified Linux or

Windows images. Each virtual machine has private virtualised hardware: a network card, disk, graphics adapter, etc. kvm is part of Linux and uses the regular Linux scheduler and memory management. This means that kvm is much smaller and simpler to use.

</software/virtual/kvm/>

#### Cooperative Linux 0.7.3:

Cooperative Linux (coLinux) is a method for optimally running Linux on Windows and other operating systems, natively. coLinux is a port of the standard Linux kernel. In other words, coLinux is the Linux kernel that's modified to run cooperatively with another operating system.

</software/virtual/colinux/>

**OpenVZ 2.6.18:** OpenVZ is a modified Linux kernel with additional support for OpenVZ Virtual Private Servers (VPS). VPSs are isolated, secure environments on a single physical server, enabling better server utilisation and ensuring that applications do not conflict. Each VPS performs and executes exactly like a stand-alone server; VPSs can be rebooted independently and have root access, users, IP addresses, memory, processes, files, applications, system libraries and configuration files.

</software/virtual/openvz/>

## For developers

**XAMPP** is a free and open source cross-platform Web server package, consisting mainly of the Apache HTTP Server, MySQL database, and interpreters for scripts written in the PHP and Perl programming languages. It is used as a development tool, to allow website designers and programmers to test their work on their own computers without any access to the Internet.

</software/developers/xampp/>

**Bob's Process Tracker** is a Linux process tracker. It lists all the processes running on your machine with details of the resources and libraries used by the process. It is useful for developers who deal with shared libraries as they can track all the shared libraries loaded at any point.

</software/developers/processtracker/>

**CodeBlocks—IDE for C++** is a free C++ IDE built to meet the most demanding needs of its users. It is designed to be very extensible and fully configurable. Built around a plug-in framework, CodeBlocks can be extended with plug-ins. Any kind of functionality can be added by installing/coding a plug-in. For instance, compiling and debugging functionality is already provided by plug-ins.

</software/developers/codeblocks/>

**KompoZer** is a complete Web authoring system that combines Web file management and easy-to-use WYSIWYG Web page editing capabilities found in Microsoft FrontPage, Adobe DreamWeaver and other popular programs.

</software/developers/kompozer/>

## X Window Programming

**Environment (xwpe)** is a programming and debugging environment similar to Borland's Turbo C environment. It works in both X and console modes. From within xwpe you can edit, compile, and debug programs.

</software/developers/xwpe/>

**MyJgui** is a graphical user interface for MySQL. You can store multiple connections that can be used simultaneously. Stored passwords are

encrypted using symmetric encryption. Underlying databases and tables are displayed in a tree structure with the connections being the first level nodes.

</software/developers/myjgui/>

## For newbies

**Avidemux** is a free video editor designed for simple cutting, filtering and encoding tasks. It supports many file types, including AVI, DVD-compatible MPEG files, MP4 and ASF, using a variety of codecs. Tasks can be automated using projects, job queues and powerful scripting capabilities.

</software/newbies/avidemux/>

**TestDisk** is a powerful free data recovery software. It was primarily designed to help recover lost partitions and/or make non-booting disks bootable again when such malfunctions are caused by faulty software, certain types of viruses or by human error (such as accidentally deleting a partition table).

</software/newbies/testdisk/>

**ISO Master** is an open source, easy-to-use, graphical CD image editor for Linux and BSD. You can use this program to extract files from an ISO, add files to an ISO, and create bootable ISOs, all in a graphical user interface. It can open ISO, NRG, and some MDF files, but can only save as ISO.

</software/newbies/isomaster/>

**Gwget** (GNOME Download Manager) is a download manager for the GNOME Desktop. By default, Gwget tries to continue any download. Gwget uses the GNOME notification area support, if available. You can close the main window and Gwget runs in the background.

</software/newbies/gwget/>

**BMP** (Beep Media Player) is a versatile and handy multi-platform media player based on the XMMS multimedia player. The BMP is mainly a port of XMMS to GTK+2 and, as such, integrates better with the look and feel of more recent versions of GNOME, Xfce, and, if using the GTK-QT theme engine, for KDE desktop environments.

</software/newbies/bmp/>

**Terminator** is an application that provides lots of terminals in a single window, saving valuable screen space otherwise wasted on window decorations that don't quite allow you to fill the screen with terminals.

</software/newbies/terminator/>

## Fun Stuff

**Gem Drop X** is a fast-paced puzzle game where it's your job to clear the screen of gems before they squash you! (You're at the bottom, they're at the top, and they keep coming at you!).

</software/funstuff/gemdropx/>

**FishWorld** is a simple and amusing point-and-click Java applet game. Fish come from both sides of the screen. The goal of the game is to let no fish through to the other side. You get points for every fish you kill (left click).

</software/funstuff/fishworld/>

**Tower Toppler**: In this game you have to help a cute little green animal switch off some kind of 'evil' mechanism.

</software/funstuff/towertoppler/>

**Step-by-Step** is a simple logic game where you have to clear all coloured tiles by stepping over them. Depending on the colour of the tile, this takes one to three steps. The game contains 99 levels and a separate level editor.

</software/funstuff/stepbystep/>

**Brickshooter** is a small puzzle game where you'll have to clear the central area of different coloured bricks. Three or more same coloured bricks that touch will vanish. You can shoot bricks into the playing field from the fringes. You can control the game with either the mouse or the keyboard.

</software/funstuff/brickshooter/>

**Ice Braker**: There's a bunch of penguins on an iceberg in Antarctica. You have been selected to catch them so they can be shipped to Finland, where they are essential to a secret plot for world domination.

</software/funstuff/icebreaker/>





This is the time when all the major GNU/Linux distributions are busy packaging their final stable releases, or squashing those bugs out of the various beta and RC releases. However, there's life beyond the 'major distributions' too. There are many not-so-well known distros that specialise in something or the other, who bring in a breath of fresh air. On that note, LFY has packed its October 2008 DVD with seven such distributions. We hope you'll give each one of them a spin and enjoy the difference. Sit back with a set of seven blank CDs, take a look at the side bar, and start burning each of these. Oh, and if you fancy virtualisation, then trying them out is even easier.

Without further ado, here's the list:

## **gOS 3 Gadgets BETA**

gOS 3 Gadgets BETA instantly launches Google Gadgets for Linux on start-up, introducing over 100,000 possible iGoogle and Google Gadgets to the desktop. Google Documents, Calendar, and Mail launch in Mozilla Prism windows to closer resemble desktop applications. The newest release of WINE 1.0 is included to now support thousands of Windows software for our advanced users. gOS 3 Gadgets BETA is based on the solid Linux distribution base of Ubuntu 8.04.1.

## **OpenGEU 8.04.1 'Luna Crescente'**

OpenGEU, previously known as Geubuntu, is a complete and fully functional OS based on the popular Linux distribution, Ubuntu. OpenGEU, a project started and designed by the Italian artist Luca D.M. (a.k.a TheDarkMaster), is targeted at any desktop, laptop or even a virtual machine. OpenGEU mixes the power and simplicity of Ubuntu and parts of the GNOME desktop with the wonder and astonishing eye-candy of Enlightenment DR17.

## **Dreamlinux Desktop Edition 3.1**

Dreamlinux is a modern and modular GNU/Linux system that can be run

directly from a CD/DVD/USB stick and optionally be installed to a HDD (IDE, SCSI, SATA, PATA and USB drive).

It comes with a selection of the best applications designed to meet most of your daily needs. Based on the venerable Debian GNU/Linux, which means it takes advantage of Debian's best features, it adds its own modern development tools, system scripts and applications. Version 3.x offers two options of desktop environments to be selected during boot time: Xfce and GNOME.

## **Linux Mint 5 'Elyssa' Main Edition (revision 1)**

Linux Mint's purpose is to produce an elegant, up-to-date and comfortable GNU/Linux desktop distribution.

Mint comes with a lot of desktop improvements that make it easier for the user to perform common tasks. There is a strong focus on making things work out of the box (Wi-Fi card drivers in the file system, multimedia support, screen resolution, etc). It is compatible with and uses Ubuntu repositories that give Linux Mint users access to a huge collection of packages and software.

## **GoblinX 2.7 Standard**

GoblinX is a bootable live CD distribution based on Slackware Linux. The primary goal for GoblinX is to create a more pleasant and functional desktop,

standardising all icons and themes to make it easy for novice users to learn about available applications.

## **Pardus Linux 2008.1 'Hyaena Hyaena'**

Pardus is a GNU/Linux distribution funded and developed by the Scientific & Technological Research Council of Turkey. Pardus has a range of unique features, such as Mudur, a start-up framework of Pardus to speed up the boot process, and PiSi, an efficient package management system with a user-friendly graphical interface. This is the KDE 4.1.1 Live CD.

## **VectorLinux 5.9 Standard Edition**

Speed, performance and stability—these are attributes that set VectorLinux apart in the crowded field of Linux distributions. The creators of VectorLinux had a single credo: keep it simple, keep it small and let the end user decide what their operating system is going to be. What has evolved from this concept is perhaps the best little Linux operating system available anywhere. The Standard Edition includes applications for every task and is specifically designed for use on older computers with slower processors and less RAM, or for those who appreciate bloat-free distributions.



# The Joy of Programming



S.G. GANESH

## Duff's Device and Some Interesting Aspects of Switch

The switch statement appears mundane; what can be special or interesting about it? In this issue, we'll explore the switch statement—you may realise that you've underestimated its value!

**C**/C++ allows only integer types for use in case statements. Why can't we use floating point numbers? Because C designers thought that it is not a good idea: checking the exact equality in floating point is not portable [ref: C99 rationale]. How about string literals? It is allowed in many languages that evolved from C, such as C#, which is a useful feature. Since switch is for integral types, a compiler can translate it to efficient code, as we will now see.

Which of the two is better: a switch statement or cascading if-else statements? Well, a switch expresses the programmer's intentions more clearly than an if-else cascade. Also, you might be surprised to know that a switch is, in general, more efficient than an equivalent if-else statement sequence! Why?

The if-else statement is flexible: it can have different conditions for each 'if' statement; also each condition can have (different) variables for comparison in the conditional expression. However, a switch statement is limited: it can have only one condition and the matching of the case statements to the condition expression is always an equality comparison; the case statements are always constant values (and not variables). Because of these reasons, the compiler can do a better job and generate efficient code. How?

A sequence of if-else statements is typically translated as a sequence of labels and jump statements (gotos). For a switch statement, a compiler generates an internal table to find the matches at runtime. Depending on the constants in the case statements, the table can be a look-up or range table. If the constants are unrelated, the comparison is usually done at the beginning and the jump is made to the specific entry in the table (i.e., a look-up table). If the constants are related and within a range (e.g., '0' to '9'), the jump can be made for each range of values (i.e., a range table). For example, a Java compiler internally compiles the switch statements into either *lookupswitch* or *tableswitch* bytecodes. So the switch is typically more efficient than if-else statements (unless the compiler is very smart, which is unlikely). The efficiency of switch statements is often exploited in different techniques and we'll now look at an unusual case.

A source of nasty bugs in C-based languages is that the case statements in the switch statement are fall-through. The 'fall-through' nature of switch is exploited in a technique called as Duff's device [Tom Duff, 'netnews', May 1984]. The following function which copies *count* number of bytes pointed by *from* to *to*:

```
send(short *to, short *from, int count){  
    do
```

```
        *to = *from++;  
    while(--count>0);  
    // this program fails if count is equal to zero.
```

and this version, compiled in a VAX C compiler, ran very slow. The reason is that the compiler translates do-while as a pair of two gotos and labels (one for each true and false case); for every condition check, a goto is executed, which makes it slow. So Tom Duff proposed another, faster version:

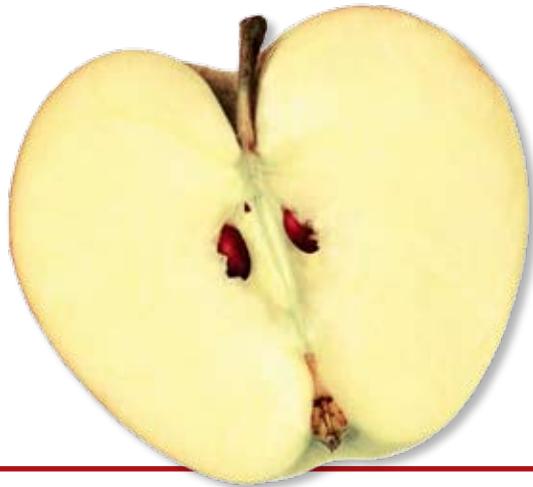
```
send(short *to, short *from, int count){  
    register n=(count+7)/8;  
    // get number of times to execute do...while loop  
    switch(count%8){  
        // go to the remaining mod value  
        case 0: do{ *to = *from++;  
        case 7: *to = *from++;  
        case 6: *to = *from++;  
        case 5: *to = *from++;  
        case 4: *to = *from++;  
        case 3: *to = *from++;  
        case 2: *to = *from++;  
        case 1: *to = *from++;  
    }while(--n>0);  
    // this loop is executed n times  
}  
// this program fails if count is equal to zero.
```

The idea is to find out the number of times the loop is to be executed in *n* and call switch to copy for modulus value. The do-while loop just ignores the case statements since they are just labels. This technique exploits the fact that case statements do not break automatically. This version ran faster than the do-while loop version (one goto for one statement) because this version has less gotos (only one goto for 8 statements) when the compiler translates it.

Even though this technique clearly exploits the fall-through nature of C switch statements, it is (fortunately) not widely used; it is good to be aware of this technique, but don't use it! 

**S.G. Ganesh** is a research engineer in Siemens (Corporate Technology). His latest book is "60 Tips on Object Oriented Programming", published by Tata McGraw-Hill. You can reach him at [sogganesh@gmail.com](mailto:sogganesh@gmail.com)

# A VOYAGE TO THE KERNEL



Part 5

## Day Four

**A** continuation of the journey of exploration, in search of all the treasures that the kernel holds!

Welcome back! We are now going to look at some more applications of shell programming. We saw the use of the *dialog* utility earlier. Now, we will learn some coding that we can incorporate into our main program so as to make it more lucrative.

To begin with, let us glance through this script that will select a random number (which is not greater than 25) and ask the user to guess the number. (If you wish to avoid the display of previous commands and messages please use the *clear* command before you start.)

```
#!/bin/sh
# Guess the random number
#Script written for A Voyage to the Kernel

biggest=25
userguess=0
totalguesses=0
varnumber=$(( $RANDOM % $biggest ))

while [ $userguess -ne $varnumber ] ; do
    echo -n "The computer has selected a number which
is less than 25. Can you guess the number? " ; read
userguess
    if [ "$userguess" -lt $varnumber ] ; then
        echo "The original number is bigger than your guessed
number!"
    elif [ "$userguess" -gt $varnumber ] ; then
        echo "The original number is smaller than your guessed
number!"
    fi
    totalguesses=$(( $totalguesses+1))
done
```

```
echo "You have guessed the number, $varnumber, in
$totalguesses guesses."
```

```
exit 0
```

Once you execute this, you'll notice that the program reads your input and churns out clues to find the correct number. It is illustrated below:

```
aasisvinayak@free-laptop:~$ ./Desktop/voyage/guessnumber.
sh
The computer has selected a number which is less than 25.
Can you guess the number? 10
The original number is bigger than your guessed number!
The computer has selected a number which is less than 25.
Can you guess the number? 15
The original number is bigger than your guessed number!
The computer has selected a number which is less than 25.
Can you guess the number? 20
The original number is bigger than your guessed number!
The computer has selected a number which is less than 25.
Can you guess the number? 24
The original number is smaller than your guessed number!
The computer has selected a number which is less than 25.
Can you guess the number? 22
You have guessed the number, 22, in 5 guesses.
```

Let's assume that you have developed an application that allows users to enter a large amount of text. Why not then incorporate a spell-check utility into your main program? Here is a code that explains what exactly happens in the shell during the process. Of course, to use the script in a program, you need to customise it, and the input data should be fed to the program. Please make sure that you have *ispell* installed before trying the script.

```
#!/bin/sh
# To check spelling a word entered
#Script written for A Voyage to the Kernel
spell="ispell -l"
for word
do
if [ -z $(echo $word | $spell) ] ; then
echo "$word -The word is spelled correctly"
else
echo "$word - The word is misspelled"
fi
done
exit 0
```

The following is a demo:

```
aasisvinayak@free-laptop:~$ ./Desktop/voyage/spellcheck.sh goat
goat -The word is spelled correctly
aasisvinayak@free-laptop:~$ ./Desktop/voyage/spellcheck.sh linux
linux - The word is misspelled
aasisvinayak@free-laptop:~$ ./Desktop/voyage/spellcheck.sh Linux
Linux -The word is spelled correctly
```

You can see from the second and third trials that for some words, it checks whether the first character is in capital letters.

We have seen the guess-number script. Now, let's discuss a guess-word script that's a little more complex.

```
#!/bin/sh
# Guess the Word (selected randomly from the list)
# Script written for A Voyage to the Kernel

blankdots="....."

selectedrandomword()
{
case $(( $$ % 8 )) in
0 ) echo "Linux" ;; 1 ) echo "GNU" ;;
2 ) echo "FSF" ;; 3 ) echo "Vlanguage" ;;

esac
}

addthegussedcharactertotheformat()
{
character=1
while [ $character -le $characters ] ; do
if [ "$(echo $word | cut -c$character)" = "$userguess" ] ;
then
before="$(( $character - 1 ))"; after="$(( $character + 1 ))"
if [ $before -gt 0 ] ; then
fbebefore="$(echo $format | cut -c$before)"
else
fbebefore=""
```

```
fi
if [ $after -gt $characters ] ; then
format="$fbebefore$userguess"
else
format="$fbebefore$userguess$(echo $format | cut -c$after-$characters)"
fi
fi
character=$(( $character + 1 ))
done

leftover=$(echo $format|sed 's/[^\.]/\./g' |wc -c|sed 's/[[[:space:]]]/\./g')
leftover=$(( $leftover - 1 ))
}

word=$((selectedrandomword))
characters=$(echo $word | wc -c | sed 's/[[[:space:]]]/\./g')
characters=$(( $characters - 1 ))
format="$(echo $blankdots | cut -c1-$characters)"
leftover=$characters ; userguessed="" ; userguesses=0;
userbadguesses=0

echo "*** Try to guess a word with $characters characters ***"

while [ $leftover -gt 0 ] ; do
echo -n "Word is: $format Guess the character next to this? ";
read userguess
userguesses=$(( $userguesses + 1 ))
if echo $userguessed | grep -i $userguess > /dev/null ; then
echo "You've already guessed that character. Try something else"
elif ! echo $word | grep -i $userguess > /dev/null ; then
echo "Sorry, the character you gussed , \"$guess\", is not in the random word selected."
userguessed="$userguessed$userguess"
userbadguesses=$(( $userbadguesses + 1 ))
else
echo "Good guess! The character $userguess is in the random word selected!"
addthegussedcharactertotheformat $userguess
fi
done
echo -n "Great! You guessed $word in $userguesses guesses"
echo " with $userbadguesses bad guesses"
exit 0
```

There is a function in the script called *addthegussedcharactertotheformat*. This function replaces the dots ('.') in the standard format with guesses. Also note that the dots "....." must be longer than the longest word in the list. And the function *selectedrandomword* will select a random word. Now, let us try executing this. (If you have observed the script carefully, you can see that there are chances of one or

two bugs emerging in it. Can you trace them?)

```
aasisvinayak@free-laptop:~$ ./Desktop/voyage/guessword.sh
** Try to guess a word with 3 characters **
Word is: ... Guess the character next to this? F
Good guess! The character F is in the random word selected!
Word is: F.F Guess the character next to this? S
Good guess! The character S is in the random word selected!
Great! You guessed FSF in 2 guesses with 0 bad guesses
aasisvinayak@free-laptop:~$ before
```

Here, when you execute the code, it asks you to enter the words. As the selected word (random) is FSF, when you enter F, two positions are filled simultaneously. Though this is not a bug, you can try removing this by adding a line that prevents the filling of more than one place simultaneously.

The code given below shows the way in which the program reacts if you enter the wrong character:

```
aasisvinayak@free-laptop:~$ ./Desktop/voyage/guessword.sh
** Try to guess a word with 3 characters **
Word is: ... Guess the letter next to this? G
Sorry, the letter you guessed , "G", is not in the random word
selected.
Word is: ... Guess the letter next to this? F
Good guess! The letter F is in the random word selected!
Word is: F.F Guess the letter next to this? S
Good guess! The letter S is in the random word selected!
Great! You guessed FSF in 3 guesses with 1 bad guesses
```

If you have not found the bug, don't worry; we are going to discuss them. Look at the following demo:

```
aasisvinayak@free-laptop:~$ /home/aasisvinayak/Desktop/guessword.
sh
** Try to guess a word with 9 characters **
Word is: ..... Guess the letter next to this? v
Good guess! The letter v is in the random word selected!
Word is: ..... Guess the letter next to this? l
Good guess! The letter l is in the random word selected!
Word is: .l..... Guess the letter next to this? a
Good guess! The letter a is in the random word selected!
Word is: .la.... Guess the letter next to this? n
Good guess! The letter n is in the random word selected!
Word is: .lan... Guess the letter next to this? g
Good guess! The letter g is in the random word selected!
Word is: .lang.ag. Guess the letter next to this? u
Good guess! The letter u is in the random word selected!
Word is: .languag. Guess the letter next to this? a
Good guess! The letter a is in the random word selected!
Word is: .languag. Guess the letter next to this? h
Sorry, the letter you guessed , "", is not in the random word
selected.
Word is: .languag. Guess the letter next to this? g
Good guess! The letter g is in the random word selected!
```

```
Word is: .languag. Guess the letter next to this? e
Good guess! The letter e is in the random word selected!
Word is: .language Guess the letter next to this? V
Good guess! The letter V is in the random word selected!
Great! You guessed Vlanguage in 11 guesses with 1 bad guesses
aasisvinayak@free-laptop:~$
```

Here the script recognises the character *v*, but it is not added. Why? As per the list, it should be in upper-case format. You can fix this by adding a statement that tells the script that both are equal.

And the following is yet another response that you may get at times:

```
** Try to guess a word with 0 characters **
Great! You guessed  in 0 guesses with 0 bad guesses
```

Why? There is an invalid decreasing range in the script. Try the `--help` command to find the resolution.

```
#!/bin/bash

echo "Set Positions"
echo '$1 = '$1
echo '$2 = '$2
echo '$3 = '$3
echo '$4 = '$4
echo '$5 = '$5
```

Now assume that you have an input box where users will enter their preferences in a particular order. You can feed that to your dynamic script (which employs `$`). And the following code is the static equivalent to achieve that:

```
aasisvinayak@free-laptop:~$ ./Desktop/voyage/arrangebyposition.sh
LFY EFY Benefit IT FFU
Set Positions
$1 = LFY
$2 = EFY
$3 = Benefit
$4 = IT
$5 = FFU
```

I am not going to illustrate this, as it is self-explanatory.

Assume that you want to add some colour to your program. For that, here is a way:

```
clear
echo -e "\033[24m Freedom"
echo -e "\033[32m Freedom"
echo -e "\033[36m Freedom"
echo -e "\033[31m Freedom"
echo -e "\033[33m Freedom"
echo -e "\033[34m Freedom"
```

```
File Edit View Terminal Tabs Help
Freedom
Freedom
Freedom
Freedom
Freedom
Freedom
Freedom
aasisvinayak@free-laptop:~$
```

Figure 1: Colours in the terminal output

```
echo -e "\033[35m Freedom"
```

Figure 1 illustrates the execution of the code.

Now if you wish to highlight the items using colours, you can try something similar to the code shown below:

```
clear
echo -e "\033[41m A Voyage to Kernel"
echo -e "\033[46m A Voyage to Kernel"
echo -e "\033[43m A Voyage to Kernel"
echo -e "\033[44m A Voyage to Kernel"
echo -e "\033[42m A Voyage to Kernel"
echo -e "\033[45m A Voyage to Kernel"
```

```
File Edit View Terminal Tabs Help
A Voyage to Kernel
aasisvinayak@free-laptop:~$
```

Figure 2: Coloured background in the terminal output

Figure 2 shows the result of the above code on execution.

Today, we have explored many ways by which you can enhance your applications. 

**By: Aasis Vinayak PG.** The author is a hacker and a free software activist who does programming in the open source domain. He is the developer and CEO of the Mozhi Search engine. His research work/publications are available at [www.aasisvinayak.com](http://www.aasisvinayak.com)

## WE TOPPED AGAIN



**BEST CERTIFIED**  
**RED HAT TRAINING PARTNER**  
**SOUTH REGION (2006-07)**



**BEST CERTIFIED**  
**RED HAT TRAINING PARTNER**  
**SOUTH REGION (2005-06)**



## RED HAT LINUX

### RHCE / RHCT / RHCSS TRACK

### RHCE EXAM ORIENTATION



#### Complete Open Source Solutions

#512, Aditya Trade Center, Ameerpet, Hyd- 38  
Tel: +91-40-66773365, Fax: +91-40-66463365  
web: <http://www.cossindia.net>

**Full Featured ThinClient System**  
Highly Efficient ThinClient Embedded OS

**Enjay Embedded ThinClient OS**  
- Linux / WinXP / WinCE

**Server Support**  
- RDP 4.0, RDP 5.x, RDP 6.x, Citrix, X11, Tarentella, NX, REXEC, VNC Client, Full Screen DOS, Novell DOS

**Local Application**  
- Open Office Suite, Email Client (Thunderbird), Web Browser (Firefox), Media Player (Audio / Video), PDF Reader, Messengers, Skype.

**Enjay Network Solutions** web : [www.enjayworld.com](http://www.enjayworld.com) | email : [info@enjayworld.com](mailto:info@enjayworld.com)  
Phone : 0260-3251732, 3241732, 3203400, 2785124 | 09898007650, 09377107650

# Open Source: A Weapon of Mass Collaboration



Does Open Source meet the four principles—openness, peering, sharing, and acting globally—proposed in *Wikinomics* as the new business imperatives?

**I**n the current age of the “prosumers” (or proactive consumers), businesses that are still marketing the traditional way are missing the beat. The Internet has placed power squarely in the hands of savvy Net users around the world. And the businesses that lead today are those that have embraced a mass collaboration model, one that enables a mutually beneficial relationship with customers by providing them with the ability to collaborate and co-innovate, giving them a voice—and choice—in a world that is shrinking as a result of uninhibited information flow.

Mass collaboration, which is built

on the foundation of open source, is a phenomenon that can no longer be ignored. An example of the booming success of mass collaboration is YouTube, the popular online video sharing site. A lesser known, but no less innovative, example exists in the Chongqing motorcycle industry in China, where the supply chain is shared among hundreds of small businesses, each focused on designing and producing a single part. The mass collaboration among the suppliers in this instance produced quality motorcycles that grew to 15 million units in just over a decade, and grabbed market share from better known Japanese and Western manufacturers.

Don Tapscott and Anthony D. Williams, two modern day thought leaders, described this business evolution as Wikinomics, defined as the “new art and science of collaboration”. In a book of the same title that is commanding the attention of enterprising business leaders throughout the world, the authors expounded on how weapons of mass collaboration is up-heaving the face of businesses today. The advent of powerful collaborative tools is giving end users across the globe the power to shape the way businesses conduct their daily operations, interact with consumers and develop new products.

The truth is, mass collaboration and open source are not new concepts. It arguably first emerged in the software world, where in 1991 a young student by the name of Linus Trovalds at the University of Helsinki released his first version of Linux, which he created as a hobby. Today, Linux represents a compelling choice to a growing number of individual consumers and enterprises that refuse to be locked in by proprietary software giants with their high licensing fees and often restrictive upgrades. There is now a proliferation of powerful open source applications, trusted by many large corporations across the globe that is supported by a strong community of developers and users, united by their common quest for widespread innovation and freedom from cloaked source codes.

The reason for the success of the open source movement is simple.

Businesses, just like consumers, want a choice of platforms, applications and service providers. They do not like being held hostage by powerful software vendors with often high-handed and exorbitant licensing policies.

However, let's examine the open source industry from a broader perspective in the light of four principles proposed in Wikinomics as the new business imperatives.

**Openness:** Open source thrives on the promise of shared knowledge, continuous innovation, and transparency. Open source software is constantly open to the injection of new ideas from its community. Source codes of applications are made available to anyone who wishes to learn it, make changes to it and improve it so that there is continuous improvement of the product. As a result, bugs in the open source software are discovered more quickly and enhancements can be made available without the need to wait for a major release. Users of open source solutions are thus not tied to a software companies' marketing time lines, where new product releases are based on perfectly timed campaigns designed to drive revenue, and not innovation as it should be.

**Mass collaboration, which is built on the foundation of open source, is a phenomenon that can no longer be ignored. An example of the booming success of mass collaboration is YouTube, the popular online video sharing site.**

**Peering:** There is no hierarchy in the open source world. Anyone in the community can contribute to the product development. Open source recognises that intelligence and skills are not limited to the top echelons of the community. This gives way to a simple paradigm—that the best software wins. Software development is no longer the privilege of elite teams, and it is no longer bundled with the latest release. Open source solutions are made available to users who can download only what they need, so they can exercise choice over the applications they wish to deploy.

**Sharing:** Open source software has grown because there is a channel to discuss ideas and issues openly which has led to the growth of a strong network and community that thrives on sharing of knowledge and expertise. In turn this gives rise to collaborative innovation and collective advancement. Anyone can be a member of this community and ideas are generated collaboratively with the best solutions chosen for deployment.

**Acting Globally:** The open source community is a global convergence of users and developers from various walks of life, brought together simply by their common interest. The diversity of the community, in itself, adds to the richness of the shared knowledge. The openness, peer support and unlimited sharing within the community epitomize globalization at its best.

That said, while most of these capability enhancing solutions are made available freely and their source codes downloadable for easy customisation, serious businesses need serious committed service providers to help them get the most out of these solutions. Businesses should seek out partners

who can package the best open source solutions with enterprise-class support with the assurance of reliability, flexibility, scalability and security.

Already, open source solutions are widely deployed in enterprise technology deployments in many corporations such as LIC, Bharti Airtel, HDFC Bank, Axis Bank, Reliance Communications, Indian Express in India, and on an APAC level there is China Telecom, AIS in Thailand, the National Bank of Australia, the University of Seoul in Korea.

The time has come to reap the true value of open source. Hail the era of unbridled innovation and the freedom from technology dictatorships. If Wikinomics is indeed the future of the global economy, then open source solutions will be the platform of choice in this new age. 

**By: Nandu Pradhan.** The author is president and MD, Red Hat India.



## Find and replace strings in any file

If we want to find a string with the name "jash", for instance, and want to replace it with the string "jassy", then we can use the *sed* command. The following is the general syntax for the *sed* command:

```
sed -i s/expression/replacement/g file.txt
```

Here, *i* is used to insert, *s* is for substitution, *expression* is what we want to find, *replacement* is what we want to replace, and *g* is used for space. So in our case, the command will be:

```
sed -i s/jash/jassy/g myfile.txt
```

If we want to take a back-up of that file before the replacement, then use the *-backup* option.

```
sed -ibackup s/jash/jassy/g myfile.txt
```

If we want to replace one string with another in the file, then we can also use the *replace* command like *sed*. Suppose we want to replace 'jash' with 'jassy', then the following is the syntax of the command:

```
replace source destination <inputfile >outputfile
```

This command will always create a new file with the new replacement.

```
replace jash jassy <file1.txt >file2.txt
```

—*Jasvendar Singh M. Chokdayat, theindianjash@gmail.com*

## Splitting and merging large files

You can split a big file of smaller parts of 100 MB each, as follows:

```
split -b 100m bigfile parts_
```

To join them in a Linux machine, use:

```
cat parts_* > bigfile
```

To join them in Windows, use:

```
copy /b parts_* bigfile
```

—*Abhilash P, webmaster@freertuxs.in*

## Know your command

If you've deleted any command, say *ping*, accidentally, or it's got corrupted for some reason unknown to you, use this tip to know the details about the loss so that you can reinstall the command.

```
# which ping  
/bin/ping
```

This is the location of the command.

```
# rpm -qf /bin/ping  
iputils-20020927-11
```

This is the RPM in which you will get the *ping* command. Install this RPM and get back your *ping*.

—*Aditi Madkaikar, aditi.madkaikar@patni.com*

## Logging in automatically

Ubuntu, by default, does not come with automatic login enabled, so whenever you start Ubuntu, you have to type in your login credentials. If you want a particular user to be logged in automatically whenever Ubuntu starts, you can do that easily by following these steps:

**Step 1:** Go to *System*—>*Administration*—>*Log-In Window*

**Step 2:** Click the *Security* tab of the new window that appears and select the *Enable Automatic login* option. Now, select the user you want to log in automatically.

—*Vijayakumar B. Patil, vizay.patil@gmail.com*

## Colourful grep

By using the `--color` option with the `grep` command, we can display the search results in different colours.

```
cat abc.txt | grep --color "xyz"
```

—Ravikumar B S, ravi\_03bs@yahoo.com

## Repeating a command sequence

If you find yourself running a sequence of commands repeatedly, `Ctrl+O` can be your friend. Press the up arrow to find the first command in the sequence, then repeatedly press `Ctrl+O` to run each command in the history.

—Ajeet Singh Raina, ajeet.singh.raina@logica.com

## What to do when the monitor goes blank

In Linux, we have found that after some time, the monitor goes into power-saving mode, i.e., it goes blank. To modify this setting, you need to run the following command:

```
xset dpms 1800 2400 3600
```

This line indicates that the monitor screen goes blank after 30 minutes (1,800 seconds), goes into power saving mode after 40 minutes (2,400 seconds) and switches off after 60 minutes (3,600 seconds).

We can change this setting according to our requirements. Also, we can use the `off` option to disable this feature.

```
xset s off
```

—Jasvendar Singh M. Chokdayat, theindianjash@gmail.com

## Shell shortcuts

While working on the shell prompt, we often type a wrong command at the beginning of a line and realise it when we've reached the end of the command. To change the starting of the command, you can press `Ctrl+A` to bring the cursor to the beginning of the line and change the first character of the command. Here's an example:

```
cd /etc/passwd
```

Note that this is a file and you cannot bring `cd` into this. So you need to go back and replace `cd` with `vi`. To do this, press `Ctrl+A`. This will take you to the start of the line and you can now replace `cd` with `vi`. Again, to go back to end of the line, press `Ctrl+E`.

—Devchandra L Meetei, dlmeetei@gmail.com

## Reclaim your Grub

Here is a process of restoring Grub, using a live CD. All you have to do is to follow the steps given below:

**Step 1:** First boot with any Linux live CD.

**Step 2:** Open a terminal and then run the following command to get the Grub command mode:

```
sudo grub
```

**Step 3:** Use the following command to find the partition that has the Grub boot loader:

```
find /boot/grub/stage1
```

**Step 4:** Note down the partition that has Grub as listed by the above command, and run the following command to make that partition your root partition:

```
root (hd?,?)
```

...where `(hd?,?)` is the partition returned by the `find` command.

**Step 5:** Next run the following command to install Grub:

```
setup (hd0)
```

This will install Grub on the Master Boot Record (MBR) of the first drive. If you want to install Grub into the boot sector of a partition instead of installing it in the MBR, specify the partition into which you want to install Grub as given below:

```
setup (hd0, 4)
```

**Step 6:** Type in the following command to quit from the Grub command line:

```
quit
```

Now reboot your computer!

—Tauqirul Haque, tauqir\_linux@yahoo.com



## Share Your Linux Recipes!

The joy of using Linux is in finding ways to get around problems—take them head on, defeat them! We invite you to share your tips and tricks with us for publication in LFY so that they can reach a wider audience. Your tips could be related to administration, programming, troubleshooting or general tweaking. Submit them at <http://www.linuxforu.com>

The sender of each published tip will get an LFY T-shirt.



NIYAM BHUSHAN

## The Harappa of the Future

They came. They coded. They conquered. And vanished.

**W**ho were they? From where did they come? Where did they go? Nothing is known about them. They existed more than 4,000 years ago. Their understanding of technology was quite advanced and unparalleled for their age. Their language: unknown. The script: undecipherable. They established trade with far away civilizations, yet fascinatingly, no civilization makes any direct reference whatsoever to the great and prosperous Indus Valley Civilization, its people, and its thoroughly advanced technology and culture. It's as if they existed in a vacuum. Leaving in their wake seals and scripts that tease us through eternity. And an enigmatic sculpture of the dancing girl of Mohenjo Daro.

### River of free knowledge?

The first city to be unearthed of this civilization was called Harappa. I recently sat through a stimulating and controversial presentation made by Dr Ravinder Bisht, the noted international archaeologist and Sanskrit-scholar. He spoke on possible insights from India's ancient Sanskrit scriptures, the Vedas, on Harappan culture. Other historians may contradict his approach. But I was quite struck with the irrefutable breadth of knowledge and technology of the Harappans, obvious from their archaeological remnants. How come they never shared their technology? Or why didn't other civilizations of the world benefit or inherit their vast and diverse knowledge? This seems almost tragic for a civilization known to exist along the banks of a perhaps mythical, perhaps real, river called Saraswati, which is also the name of the Hindu Goddess of Knowledge.

### Ravi Indiana Jones Kant

When it comes to history's burning questions such as these, I turn to India's muft and mukt incarnation of Indiana Jones for answers. His name is Ravi Kant. This desi Jones researched and taught history for a number of years at Delhi University. He currently works with Sarai, which is a famous and FOSS-friendly project of the Centre for the Study of Developing Societies (CSDS). Incidentally, the word 'Sarai' literally means an enclosed space in a city or a tavern. Ravi Kant is also one of the leading and yet unsung heroes of India's Linux and FOSS community. Just the sort of guy to ask about the knowledge of ancient India. Alas! Ravi just sums it up in

two sentences: "I don't think there was anything inherent in the Harappan civilization that disallowed the continuity of knowledge. I think it is a comment on us that we have not yet been able to decipher their script."

### Digital Harappa

Reminds me of how inspired I felt at the end of Dr Bisht's presentation. I imagined a time 4,000 years into the future. We stand to become the Harappan Civilization of the future, if we do not embrace the culture of sharing today. Future civilizations may discover relics of our currently revolutionary digital technologies. But if these are not muft and mukt, they may just become puzzling enigmas of the archeology of the digital. If we do not share and contribute our knowledge to a copyleft world today, the future may not even see any reference to us in the other civilizations of today. The shining country in today's

BRIC alliance, may just become another brick in the wall of oblivion. Just like the wealth and the prosperity of the Harappan civilization vanished, so would ours. Beyond the crushing deadlines of quarterly balance-sheets,

history will eventually reveal to us how wealth is inexorably linked to the knowledge-culture of a civilization.

In that light, the Sarai's ground-breaking work in FOSS-based Indic localisation is far-reaching. Likewise, their FOSS scholarships and various initiatives are commendable. But to survive, endure and prosper, it is not one lone Sarai, but the entire Indian sub-continent that has to ignite with the culture of sharing knowledge digitally, within the value-systems of FOSS.

Seeking further inspiration, I turn to Osho and am startled to find it in the opening sentence of his introduction to the circa 5,000 year-old Indian meditation scripture, the Vigyan Bhairav Tantra: "Truth is always here." Now that alone is truly worth knowing. 

*Inspired by the vision of Osho. Copyright September 2008: Niyam Bhushan. freedomyugs\_at\_gmail\_dotcom. First published in LinuxForYou magazine. Verbatim copying, publishing and distribution of this article is encouraged in any language and medium, so long as this copyright notice is preserved. In Hindi, 'muft' means 'free-of-cost', and 'mukt' means 'with freedom'.*

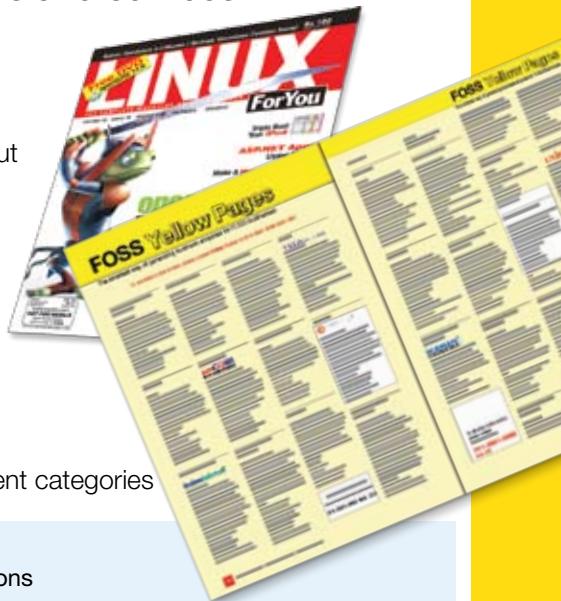
GUEST COLUMN

# FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

## HIGHLIGHTS

- A cost-effective marketing tool
- A user-friendly format for customers to contact you
- A dedicated section with yellow back-ground, and hence will stand out
- Reaches to tech-savvy IT implementers and software developers
- 80% of LFY readers are either decision influencers or decision takers
- Discounts for listing under multiple categories
- Discounts for booking multiple issues



## FEATURES

- Listing is categorised on the basis of products and services
- Complete contact details plus 30-word description of organisation
- Option to print the LOGO of the organisation too (extra cost)
- Option to change the organisation description for listings under different categories

## TARIFF

### Category Listing

ONE Category .....	Rs 2,000
TWO Categories.....	Rs 3,500
THREE Categories.....	Rs 4,750
ADDITIONAL Category .....	Rs 1,000

### Value-add Options

LOGO-plus-Entry.....	Rs 500
Highlight Entry (white background).....	Rs 1,000
Per EXTRA word (beyond 30 words).....	Rs 50

## KEY POINTS

- Above rates are per-category basis.
- Above rates are charges for publishing in a single issue of LFY.
- Max. No. of Words for Organisation Description: 30 words.

## TERMS & CONDITIONS

- Fill the form (below).
- You can use multiple copies of the form for multiple listings under different categories.
- Payment to be received along with booking.

 Tear & Send

 Tear & Send

## ORDER FORM

Organisation Name (70 characters): \_\_\_\_\_

Description (30 words): \_\_\_\_\_

Email: \_\_\_\_\_ Website: \_\_\_\_\_

STD Code: \_\_\_\_\_ Phone: \_\_\_\_\_ Mobile: \_\_\_\_\_

Address (will not be published): \_\_\_\_\_

City/Town: \_\_\_\_\_ Pin-code: \_\_\_\_\_

## CATEGORIES

<input type="checkbox"/> CONSULTANTS	<input type="checkbox"/> HIGH PERFORMANCE COMPUTING	<input type="checkbox"/> SOFTWARE DEVELOPMENT
<input type="checkbox"/> CONSULTANT (FIRM)	<input type="checkbox"/> IT INFRASTRUCTURE SOLUTIONS	<input type="checkbox"/> TRAINING FOR PROFESSIONALS
<input type="checkbox"/> EMBEDDED SOLUTIONS	<input type="checkbox"/> LINUX-BASED WEB-HOSTING	<input type="checkbox"/> TRAINING FOR CORPORATE
<input type="checkbox"/> ENTERPRISE COMMUNICATION SOLUTIONS	<input type="checkbox"/> MOBILE SOLUTIONS	<input type="checkbox"/> THIN CLIENT SOLUTIONS

Please find enclosed a sum of Rs. \_\_\_\_\_ by DD/ MO//crossed cheque\* bearing the No. \_\_\_\_\_ dt. \_\_\_\_\_ in favour of

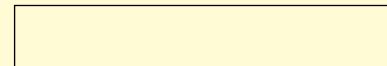
EFY Enterprises Pvt Ltd, payable at Delhi. (\*Please add Rs. 50 on non-metro cheque) towards the cost of \_\_\_\_\_ FOSS Yellow Pages advertisement(s)

or charge my credit card  VISA  Master Card Please charge Rs. \_\_\_\_\_

against my credit card No. 

C V V No. \_\_\_\_\_ (Mandatory)

Date of Birth \_\_\_\_\_ / \_\_\_\_\_ / \_\_\_\_\_ (dd/mm/yy) Card Expiry Date \_\_\_\_\_ / \_\_\_\_\_ (mm/yy)



Signature (as on the card)



EFY Enterprises Pvt Ltd., D-87/1, Okhla Industrial Area, Phase 1, New Delhi 110 020  
Ph: 011-26810601-03, Fax: 011-26817565, Email: info@efyindia.com; Website: www.efyindia.com

**To Book Your Listing, Call:** Somaiah (B'llore: 09986075717) Suraj (Delhi: 09350475845)

# FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

To advertise in this section, please contact:

Somaiah (Bangalore) **09986075717**, Suraj (Delhi) **09350475845**

## Consultant (Firm)



### Mentor Infotech Solutions

Trusted name for implementation of Enterprise Workflow Solutions. Expertise in Java based portal implementation, End-to-End IT Project Consultancy, Networks and Security management.

**Lucknow-226010**

Phone: 0522-3291999, 4064430

Email: [mentor@sancharnet.in](mailto:mentor@sancharnet.in)

Web: [www.mentorinfotech.in](http://www.mentorinfotech.in)

### TechnoInfotech™

A company focussed on Enterprise Solution using opensource software. Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

**Mumbai-400001**

Phone: 022-66338900; Extn. 324

Email: [sales@technoinfotech.com](mailto:sales@technoinfotech.com)

Web: [www.technoinfotech.com](http://www.technoinfotech.com)

## Enterprise Communication Solutions

### Netcore Solutions Pvt. Ltd.

Emeric CleanMail is an Internet-based filtering service to monitor and filter e-mail traffic to protect against virus attacks, spam mails and wasted bandwidth. It offers triple-level virus scanning with extensive graphical reports on mail usage patterns. Emeric MailArchiva - Mail Archiving solutions for corporates which can get integrated with all mailing platform.

**Mumbai-400013**

Phone: 022-66628000

Mobile: 09322985222

Email: [kalpit@netcore.co.in](mailto:kalpit@netcore.co.in)

Web: [www.netcore.co.in](http://www.netcore.co.in)

To advertise in this section,

please contact

Somaiah (B'lore):

09986075717

Suraj (Delhi): 09350475845) on

**011-2681-0602 Extn. 222**



### Red Hat India Pvt. Ltd.

Red Hat is the world's leading open source solutions provider. Red Hat provides high-quality, affordable technology with its operating system platform, Red Hat Enterprise Linux, together with applications, management and Services Oriented Architecture (SOA) solutions, including JBoss Enterprise Middleware. Red Hat also offers support, training and consulting services to its customers worldwide.

**Mumbai-400076**

Phone: 022-39878888

Email: [marketing-in@redhat.com](mailto:marketing-in@redhat.com)

Web: [www.redhat.in](http://www.redhat.in)

### TechnoInfotech™

A company focussed on Enterprise Solution using opensource software. Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

**Mumbai-400001**

Phone: 022-66338900; Extn. 324

Email: [sales@technoinfotech.com](mailto:sales@technoinfotech.com)

Web: [www.technoinfotech.com](http://www.technoinfotech.com)

## High Performance Computing

### Netcore Solutions Pvt. Ltd.

Emeric MailServ offers an integrated solution for email, IM, proxy, global address book, firewall, VPN, bandwidth management, anti-virus, anti-spam and content filtering. It has an easy-to-use remote management dashboard.

**Mumbai-400013**

Phone: 022-66628000

Mobile: 09322985222

Email: [kalpit@netcore.co.in](mailto:kalpit@netcore.co.in)

Web: [www.netcore.co.in](http://www.netcore.co.in)

## IT Infrastructure Solutions

### BakBone Software Inc.

BakBone Software Inc. delivers complexity-reducing data protection technologies, including award-winning Linux solutions; proven Solaris products; and application-focused Windows offerings that reliably protect MS SQL, Oracle,

Exchange, MySQL and other business critical applications.

**New Delhi-110048**

Tel: 011-42235156

Email: [ashish.gupta@bakbone.com](mailto:ashish.gupta@bakbone.com)

Web: [www.bakbone.com](http://www.bakbone.com)

### Keen & Able Computers Pvt. Ltd.

Open Source Solutions Provider. Red Hat Ready Business Partner. Mail Servers/Anti-spam/GUI interface/Encryption, Clustering & Load Balancing - SAP/Oracle/Web/Thin Clients, Network and Host Monitoring, Security Consulting, Solutions, Staffing and Support.

**New Delhi-110019**

Tel: 011-30880046, 30880047

Mobile: 09810477448, 09891074905

Email: [info@keenable.com](mailto:info@keenable.com)

Web: [www.keenable.com](http://www.keenable.com)

### TechnoInfotech™

A company focussed on Enterprise Solution using opensource software.

Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

**Mumbai-400001**

Phone: 022-66338900; Extn. 324

Email: [sales@technoinfotech.com](mailto:sales@technoinfotech.com)

Web: [www.technoinfotech.com](http://www.technoinfotech.com)

## Software Development

### Carizen Software (P) Ltd.

Carizen's flagship product is Rainmail Intranet Server, a complete integrated software product consisting modules like mail sever, proxy server, gateway anti-virus scanner, anti-spam, groupware, bandwidth aggregator & manager, firewall, chat server and fax server. Infrastructure.

**Chennai-600028**

Phone: 044-24958222, 8228, 9296

Email: [info@carizen.com](mailto:info@carizen.com)

Web: [www.carizen.com](http://www.carizen.com)

## CATEGORIES FOR FOSS YELLOW PAGES

- Consultants
- Consultant (Firm)
- Embedded Solutions
- Enterprise Communication Solutions
- High Performance Computing
- IT Infrastructure Solutions
- Linux-based Web-hosting
- Mobile Solutions
- Software Development
- Training for Professionals
- Training for Corporate
- Thin Client Solutions

To advertise in this section, please contact

Somaiah (Bangalore) **09986075717**

Suraj (Delhi) **09350475845**

## Linux Based Web Hosting

### TechnoInfotech™

A company focussed on Enterprise Solution using opensource software. Key Solutions:

- Enterprise Email Solution
- Internet Security and Access Control
- Managed Services for Email Infrastructure.

Mumbai-400001

Phone: 022-66338900; Extn. 324

Email: sales@technoinfotech.com

Web: www.technoinfotech.com

## Training for Corporate

### Complete Open Source Solutions

RHCT, RHCE and RHCSA training.

**Hyderabad-500038**

Phone: 040-66773365, 9849742065

Email: nayak.sujeet@gmail.com

Web: www.cossindia.com

### Lynus Academy Pvt. Ltd.

India's premier Linux and OSS training institute.

Chennai-600101

Phone: 044-42171278, 9840880558

Email: contactus@lynusacademy.com

Web: www.lynusacademy.com

### Linux Learning Centre Private Limited

Pioneers in training on Linux technologies.

Bangalore-560019

Phone: 080-22428538, 26600839

Email: info@linuxlearningcentre.com

Web: www.linuxlearningcentre.com

### Netweb Technologies

Simplified and scalable storage solutions.

Bangalore-560001

Phone: 080-41146565, 32719516

Email: info@netwebindia.com

Web: www.netwebindia.com

## Thin Client Solutions

### Enjay Network Solutions

Gujarat based ThinClient Solution Provider. Providing Small Size ThinClient PCs & a Full Featured ThinClient OS to perfectly suite needs of different working environment. Active Dealer Channel all over India. Gujarat-396105

Tel.: 0260-3203400, 3241732, 3251732,

Mobile: 09377107650, 09898007650

Email: info@enjayworld.com

Web: www.enjayworld.com

To advertise in this section, please contact

Somaiah (Bangalore) 09986075717

Suraj (Delhi) 09350475845

## FOSS Yellow Pages

## New Horizons India Ltd.

New Horizons India Ltd., a joint venture of New Horizons Worldwide, Inc. (NASDAQ: NEWH) and the Shriram group, is an Indian company operational since 2002 with a global footprint engaged in the business of knowledge delivery through acquiring, creating, developing, managing, lending and licensing knowledge in the areas of IT, Applied Learning, Technology Services and Supplementary Education. The company has pan India presence with 15 offices and employs 750 people.

New Delhi-110003

Tel: 011-43612400

Email: info@nhindia.com

Web: www.nhindia.com

## Training for Professional

### FOSTERing Linux

Linux & Open Source Training Institute, All trainings provided by experienced experts & System Administrators only, RHCE, RHCSA, (Red Hat Training & Examination Partners), PHP, Perl, OpenOffice, Clustering, Mail Servers, Bridging the GAP by providing: Quality training (corporate & individual), Quality Manpower, Staffing and Support & 100% Placement Assistance.

Gurgaon-122001

Tel: 0124-4268187, 4080880

Mobile: 09350640169, 09818478555

Email: info@fosteringlinux.com

Web: www.fl.keenable.com

# FOSS Yellow Pages

The best place for you to buy and sell FOSS products and services

# LINUX For You

Calling...

Share your  
knowledge

Inviting FOSS experts to write articles  
on their area of interest

LFY covers a myriad of topics—network management, software development, embedded systems, community issues, and even hands-on guide for newbies. If you've got an interesting topic, let us know. Thanks to the launch of [www.openLTis.com](http://www.openLTis.com) (aka [linuxforu.com](http://linuxforu.com)) we are now trying to extend our content portfolio related to Linux & Open Source.

To know more on how to become an LFY author, contact us at

[lfyedit@efyindia.com](mailto:lfyedit@efyindia.com)

### Few topics that top our list:

- ◆ Tips 'n' Tricks for software developers or IT implementers
- ◆ Cool tweaks for FOSS enthusiasts
- ◆ FOSS on mobile
- ◆ Virtualisation (Implementation)
- ◆ OpenJDK or Java on Linux
- ◆ OpenSolaris (software development)
- ◆ How can I do 'that' on Linux
- ◆ Reviews of latest open source projects & tools



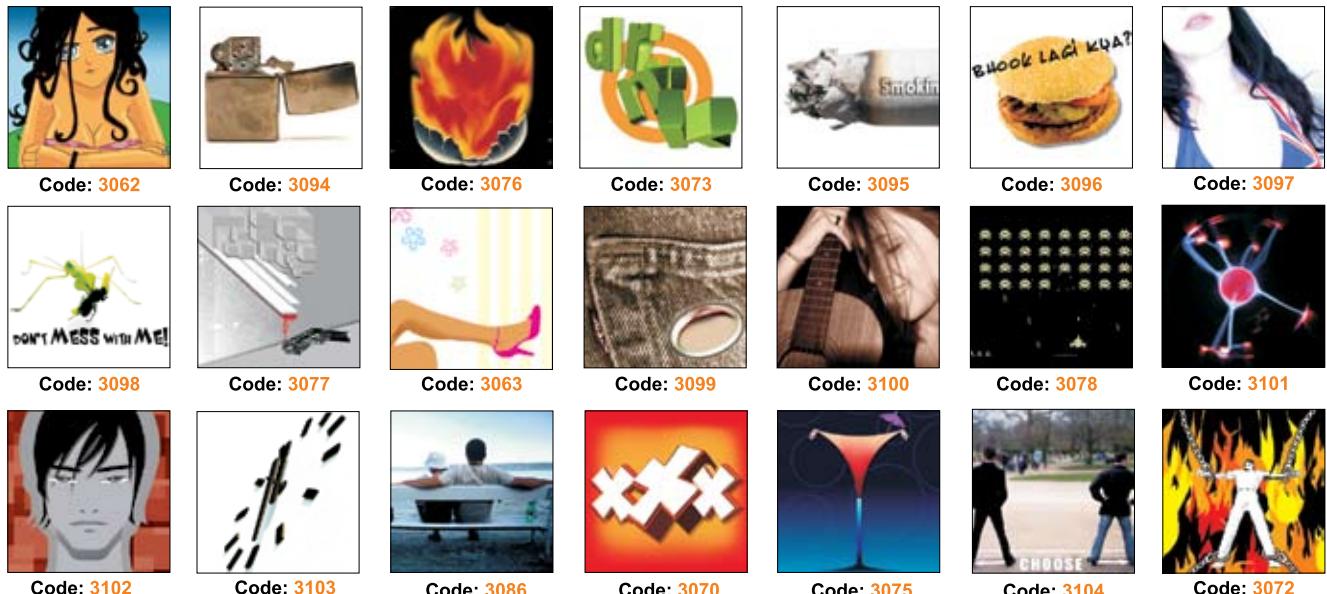
**Indipop Special Tones (Mono & Polytones)**

Song name	Movie/Album name	Code
Ehsaas	Doorie	3430
Mahi Ve	Doorie	3431
Yakeen (Club Mix)	Doorie	3432
Sajnaa	Sajnaa	3433
Nach Le Soniye Tu	Shaels Aitbaar	3434

**Hot Ringtones (Mono & Polytones)**

Song name	Movie/Album name	Code
Rutho Naa	With Luv Tumhaara	3435
Kyun Ho Khafa	With Luv Tumhaara	3436
Dilbar Mera	Anwar	3437
Javeda - Tose Naina Lagey	Anwar	3438
Maula Mere Maula	Anwar	3439

**Wallpapers**



**Picture Messages**



**Download Instructions**

**Wallpapers:** SMS WALL <code> to 53636. (GPRS services required)

**Picture messages:** SMS PM<code> to 53636. (Handset should be picture message compatible)

**Monotones:** For Nokia/ Samsung/ LG phones: SMS RT <code> to 53636  
Others SMS ORT <code> to 53636

**Polytones:** SMS POLY <code> to 53636 (GPRS services required)

VAS SMS rates apply. \*BPL Wallpapers & Polytones available in Mumbai only.

\*\*Wallpapers & Polytones not available on Reliance GSM, IDEA & BSNL.

\*\*\*Only picture messages available on Spice Punjab

**Content download cost (in Rs.)**

Operator	Monotone	Polytone	Pic. Msg.	Wallpaper
Airtel	10/-	13/-	3/-	13/-
BPL*	7/-	15/-	3/-	10/-
Hutch	10/-	13/-	3/-	13/-
MTNL-Mumbai	7/-	9/-	3/-	9/-
MTNL-Delhi	8/-	10/-	3/-	10/-
Reliance GSM**	7/-	NA	3/-	NA
BSNL**	7/-	NA	3/-	NA
IDEA**	10/-	NA	8/-	NA
Spice (Punjab)***	NA	NA	8/-	NA



Future is Open...

# **NOKIA**

## Connecting People

The Linux-based Maemo platform is the software that powers mobile devices such as the Nokia N810 Internet Tablet. To make the Maemo platform even more powerful, we at the Maemo SW entity in Nokia Devices R&D are now looking for people to take it to the next level as part of the applications development team. We aim to extend the platform's capabilities as software for innovative communication and Internet devices.

Being part of Nokia's Devices R&D business group gives us a world-class environment to use Linux and open source technologies in new and interesting ways.

---

### **Senior Developer**

You have at least 5 years experience in SW development for embedded devices on Linux including C/C++ and Qt/Qtopia skills. You know how SW development is conducted in open source projects and are familiar with the legal and licensing issues related to them as well as the Agile-Scrum SW development model.

---

In addition, you are a solid performer with a strong drive to meet the targets. You enjoy learning and expanding yourself.

---

### **Project Manager**

You have total industry experience of 8 years with at least last 4 years in Project Management of SW development projects. A successful candidate should have experience in SW development as well as in PMI-style and agile PM. You are responsible for the planning, estimation, delivery and quality assurance of components in our application projects.

Mail your resumes to [amit.jaiswal@nokia.com](mailto:amit.jaiswal@nokia.com)



Welcome to Microsoft Virtualization.

**Microsoft Virtualization breaks down barriers to creating the Virtual Enterprise. With end-to-end solutions, not only can you manage your technology infrastructure in an easy, smart and flexible manner, but also accelerate IT capabilities while reducing costs.**

© 2008 Microsoft Corporation. All rights reserved. Microsoft and 'Your potential. Our Passion' are all registered trademarks of Microsoft Corporation in the United States and/or other countries.

**Microsoft® Virtualization**